



(11) **EP 4 276 608 A2**

(12) **EUROPEAN PATENT APPLICATION**

- (43) Date of publication: **15.11.2023 Bulletin 2023/46**
- (51) International Patent Classification (IPC): **G06F 9/30 (2018.01)**
- (21) Application number: **23195872.9**
- (52) Cooperative Patent Classification (CPC): **G06F 9/30014; G06F 9/30025; G06F 9/30036**
- (22) Date of filing: **14.09.2021**

(84) Designated Contracting States:
AL AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HR HU IE IS IT LI LT LU LV MC MK MT NL NO PL PT RO RS SE SI SK SM TR

(30) Priority: **26.12.2020 US 202017134373**

(62) Document number(s) of the earlier application(s) in accordance with Art. 76 EPC:
21196474.7 / 4 020 169

(71) Applicant: **Intel Corporation**
Santa Clara, CA 95054 (US)

- (72) Inventors:
- **Mellempudi, Naveen**
560035 Bangalore (IN)
 - **Heinecke, Alexander F.**
San Jose, CA 95134 (US)
 - **Valentine, Robert**
36054 Kiryat Tivon HA (IL)

- **Charney, Mark J.**
Lexington, MA 02421 (US)
- **Hughes, Christopher J.**
Santa Clara, CA 95051 (US)
- **Georganas, Evangelos**
San Jose, CA 95129 (US)
- **Sperber, Zeev**
3092832 Zichron Yackov (IL)
- **Gradstein, Amit**
3052316 Binyamina HA (IL)
- **Rubanovich, Simon**
34792 Haifa HA (IL)

(74) Representative: **Samson & Partner Patentanwälte mbB**
Widenmayerstraße 6
80538 München (DE)

Remarks:

This application was filed on 07-09-2023 as a divisional application to the application mentioned under INID code 62.

(54) **APPARATUSES, METHODS, AND SYSTEMS FOR 8-BIT FLOATING-POINT MATRIX DOT PRODUCT INSTRUCTIONS**

(57) Systems, methods, and apparatuses relating to 8-bit floating-point matrix dot product instructions are described. For example, a processing unit comprises circuitry to perform operations corresponding to an instruction, the instruction to specify a first matrix having M rows by 4*K columns of 8-bit floating-point data elements, a second matrix having 4*K rows by N columns of 8-bit floating-point data elements, and a third matrix having M rows by N columns of 32-bit single precision floating-point data elements. The operations includes to, for each row m of the M rows of the first matrix, and for each column n of the N columns of the second matrix: convert 4*K 8-bit floating-point data elements of the row m of the first matrix to 4*K corresponding higher precision floating-point data elements having a higher precision than an 8-bit floating-point data element, and convert 4*K 8-bit floating-point data elements of the column n of the second matrix to 4*K corresponding higher precision floating-point data elements having a higher precision than the 8-bit floating-point data element; multiply the 4*K higher precision floating-point data elements corresponding to the row m of the first matrix with correspond-

ing ones of the 4*K higher precision floating-point data elements corresponding to the column n of the second matrix to generate 4*K products; accumulate the 4*K products with a 32-bit single precision floating-point data element corresponding to a row m of the M rows, and a column n of the N columns, of the third matrix, to generate a result 32-bit single precision floating-point data element; and store the result 32-bit single precision floating-point data element at the row m and the column n of the third matrix.




FIG. 1A

EP 4 276 608 A2

Description**TECHNICAL FIELD**

5 **[0001]** The disclosure relates generally to computer processor architecture, and, more specifically, to systems and methods for performing 8-bit floating-point matrix dot product instructions.

BACKGROUND

10 **[0002]** Matrices are increasingly important in many computing tasks such as machine learning and other bulk data processing. Deep Learning is a class of machine learning algorithms. Deep learning architectures, such as deep neural networks, have been applied to fields including computer vision, speech recognition, natural language processing, audio recognition, social network filtering, machine translation, bioinformatics and drug design.

15 **[0003]** Inference and training, two tools used for deep learning, are tending towards low precision arithmetic. Maximizing throughput of deep learning algorithms and computations may assist in meeting the needs of deep learning processors, for example, those performing deep learning in a data center.

[0004] Matrix-matrix multiplication (a.k.a., GEMM or General Matrix Multiplication) is a common compute-heavy operation on modern processors. Special hardware for matrix multiplication (e.g., GEMM) is a good option for improving the peak compute (and energy efficiency) of certain applications, such as deep learning.

20 **[0005]** Some of these applications, including deep learning, can operate on input data elements with relatively few bits without losing accuracy, as long as the output elements have enough bits (i.e., more than the inputs).

BRIEF DESCRIPTION OF THE DRAWINGS

25 **[0006]** The present disclosure is illustrated by way of example and not limitation in the figures of the accompanying drawings, in which like references indicate similar elements and in which:

Figure 1A illustrates an embodiment of configured tiles;

Figure 1B illustrates an embodiment of configured tiles;

30 **Figure 2** illustrates several examples of matrix storage;

Figure 3 illustrates an embodiment of a system utilizing a matrix (tile) operations accelerator;

Figures 4 and 5 show different embodiments of how memory is shared using a matrix operations accelerator;

Figure 6 illustrates an embodiment of matrix multiply accumulate operation using tiles ("TMMA");

35 **Figure 7** illustrates an embodiment of a subset of the execution of an iteration of a chained fused multiply accumulate instruction;

Figure 8 illustrates an embodiment of a subset of the execution of an iteration of a chained fused multiply accumulate instruction;

Figure 9 illustrates an embodiment of a subset of the execution of an iteration of a chained fused multiply accumulate instruction;

40 **Figure 10** illustrates an embodiment of a subset of the execution of an iteration of chained fused multiply accumulate instruction;

Figure 11 illustrates power-of-two sized SIMD implementations wherein the accumulators use input sizes that are larger than the inputs to the multipliers according to an embodiment;

Figure 12 illustrates an embodiment of a system utilizing matrix operations circuitry;

45 **Figure 13** illustrates an embodiment of a processor core pipeline supporting matrix operations using tiles;

Figure 14 illustrates an embodiment of a processor core pipeline supporting matrix operations using tiles;

Figure 15 illustrates an example of a matrix expressed in row major format and column major format;

Figure 16 illustrates an example of usage of matrices (tiles);

Figure 17 illustrates an embodiment a method of usage of matrices (tiles);

50 **Figure 18** illustrates support for configuration of the usage of tiles according to an embodiment;

Figure 19 illustrates an embodiment of a description of the matrices (tiles) to be supported;

Figures 20(A)-(D) illustrate examples of register(s);

Figure 21A is a block diagram illustrating use of a TDPBF8PS instruction to accelerate matrix multiplication, according to some embodiments;

55 **Figure 21B** is a block diagram illustrating example execution circuitry to execute a TDPBF8PS instruction, according to some embodiments;

Figure 22A is pseudocode illustrating execution of a TDPBF8PS instruction according to some embodiments;

Figure 22B is pseudocode illustrating helper functions for use by the pseudocode of Figure 22A, according to some

embodiments;

Figure 23 illustrates an embodiment of a processor executing a flow to process a TDPBF8PS instruction;

Figure 24 is a block diagram illustrating a format of a TDPBF8PS instruction according to some embodiments;

Figures 25A-25B are block diagrams illustrating a generic vector friendly instruction format and instruction templates thereof according to embodiments;

Figure 25A is a block diagram illustrating a generic vector friendly instruction format and class A instruction templates thereof according to embodiments;

Figure 25B is a block diagram illustrating the generic vector friendly instruction format and class B instruction templates thereof according to embodiments;

Figure 26A is a block diagram illustrating an exemplary specific vector friendly instruction format according to embodiments;

Figure 26B is a block diagram illustrating the fields of the specific vector friendly instruction format that make up the full opcode field according to one embodiment;

Figure 26C is a block diagram illustrating the fields of the specific vector friendly instruction format that make up the register index field according to one embodiment;

Figure 26D is a block diagram illustrating the fields of the specific vector friendly instruction format that make up the augmentation operation field according to one embodiment;

Figure 27 is a block diagram of a register architecture according to one embodiment;

Figure 28A is a block diagram illustrating both an exemplary in-order pipeline and an exemplary register renaming, out-of-order issue/execution pipeline according to embodiments;

Figure 28B is a block diagram illustrating both an exemplary embodiment of an in-order architecture core and an exemplary register renaming, out-of-order issue/execution architecture core to be included in a processor according to embodiments;

Figures 29A-B illustrate a block diagram of a more specific exemplary in-order core architecture, which core would be one of several logic blocks (including other cores of the same type and/or different types) in a chip;

Figure 29A is a block diagram of a single processor core, along with its connection to the on-die interconnect network and with its local subset of the Level 2 (L2) cache, according to embodiments;

Figure 29B is an expanded view of part of the processor core in **Figure 29A** according to embodiments;

Figure 30 is a block diagram of a processor that may have more than one core, may have an integrated memory controller, and may have integrated graphics according to embodiments;

Figures 31-34 are block diagrams of exemplary computer architectures;

Figure 31 shown a block diagram of a system in accordance with one embodiment of the present disclosure;

Figure 32 is a block diagram of a first more specific exemplary system in accordance with an embodiment of the present disclosure;

Figure 33 is a block diagram of a second more specific exemplary system in accordance with an embodiment of the present disclosure;

Figure 34 is a block diagram of a System-on-a-Chip (SoC) in accordance with an embodiment of the present disclosure; and

Figure 35 is a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction set according to embodiments.

DETAILED DESCRIPTION

[0007] In the following description, numerous specific details are set forth. However, it is understood that embodiments may be practiced without these specific details. In other instances, well-known circuits, structures and techniques have not been shown in detail in order not to obscure the understanding of this description.

[0008] References in the specification to "one embodiment," "an embodiment," "an example embodiment," etc., indicate that the embodiment described may include a particular feature, structure, or characteristic, but every embodiment may not necessarily include the particular feature, structure, or characteristic. Moreover, such phrases are not necessarily referring to the same embodiment. Further, when a particular feature, structure, or characteristic is described in connection with an embodiment, it is submitted that it is within the knowledge of one skilled in the art to affect such feature, structure, or characteristic in connection with other embodiments whether or not explicitly described.

[0009] In many mainstream processors, handling matrices is a difficult and/or instruction intensive task. For example, rows of a matrix could be put into a plurality of packed data (e.g., SIMD or vector) registers and then operated on individually. For example, an add two 8x2 matrices may require a load or gather into four packed data registers depending upon data sizes. Then a first add of packed data registers corresponding to a first row from each matrix is performed and a second add of packed data registers corresponding to a second row from each matrix is performed. Then the resulting packed data registers are scattered back to memory. While for small matrices this scenario may be acceptable,

it is often not acceptable with larger matrices.

DISCUSSION

5 **[0010]** Described herein are mechanisms to support matrix operations in computer hardware such as central processing units (CPUs), graphic processing units (GPUs), and accelerators. The matrix operations utilize 2-dimensional (2-D) data structures representing one or more packed regions of memory such as registers. Throughout this description, these 2-D data structures are referred to as tiles. Note that a matrix may be smaller than a tile (use less than all of a tile) or utilize a plurality of tiles (the matrix is larger than the size of any one tile). Throughout the description, matrix (tile) language is used to indicate operations performed using tiles that impact a matrix; whether or not that matrix is larger than any one tile is not typically relevant.

10 **[0011]** Each tile may be acted upon by different operations such as those that are detailed herein and include, but are not limited to: matrix (tile) multiplication, tile add, tile subtract, tile diagonal, tile zero, tile transform, tile dot product, tile broadcast, tile row broadcast, tile column broadcast, tile multiplication, tile multiplication and accumulation, tile move, etc. Additionally, support for operators such as the use of a scale and/or bias may be used with these operations or in support of non-numeric applications in the future, for instance, OpenCL "local memory," data compression/decompression, etc. Also described herein are instructions for performing matrix (tile) 8-bit floating-point tile dot product (TDPBF8PS) instructions.

15 **[0012]** Portions of storage (such as memory (non-volatile and volatile), registers, cache, etc.) are arranged into tiles of different horizontal and vertical dimensions. For example, a tile may have horizontal dimension of 4 (e.g., four rows of a matrix) and a vertical dimension of 8 (e.g., 8 columns of the matrix). Typically, the horizontal dimension is related to element sizes (e.g., 2-, 4-, 8-, 16-, 32-, 64-, 128-bit, etc.). Multiple datatypes (single precision floating-point, double precision floating-point, integer, etc.) may be supported.

25 EXEMPLARY USAGE OF CONFIGURED TILES

[0013] In some embodiments, tile parameters can be configured. For example, a given tile may be configured to provide tile options. Exemplary tile options include but are not limited to: a number of rows of the tile, a number of columns of the tile, whether the tile is VALID, and/or whether the tile consists of a PAIR of equal-sized tiles.

30 **[0014]** **Figure 1A** illustrates an embodiment of configured tiles. As shown, 4 kB of application memory 102 have stored thereon 4 1kB tiles, tile t0 104, tile t1 106, tile t2 108, and tile t3 110. In this example, the 4 tiles do not consist of pairs, and each have elements arranged in rows and columns. Tile t0 104 and tile t1 106 have K rows and N columns of 4-byte elements (e.g., single precision data), where K equals 8 and N=32. Tile t2 108 and tile t3 110 have K rows and N/2 columns of 8-byte elements (e.g., double precision data). As the double precision operands are twice the width of single precision, this configuration is consistent with a palette, used to provide tile options, supplying at least 4 names with total storage of at least 4 kB. In operation, the tiles can be loaded from and stored to memory using load and store operations. Depending upon the instruction encoding scheme used, the amount of available application memory, as well as the size, number, and configuration of available tiles varies.

35 **[0015]** **Figure 1B** illustrates an embodiment of configured tiles. As shown, 4 kB of application memory 122 have stored thereon 2 pairs of 1kB-tiles, the first pair being tile t4L 124 and tile t4R 126, and the second pair being tile t5L 128 and tile t5R 130. As shown the pairs of tiles are divided into a left tile and a right tile. In other embodiments, the pair of tiles are divided into an even tile and an odd tile. In this example, the 4 tiles each have elements arranged in rows and columns. Tile t4L 124 and tile t4R 126 have K rows and N columns of 4-byte elements (e.g., single precision floating-point data), where K equals 8 and N equals 32. Tile t5L 128 and tile t5R 130 have K rows and N/2 columns of 8-byte elements (e.g., double precision floating-point data). As the double precision operands are twice the width of single precision, this configuration is consistent with a palette, used to provide tile options, supplying at least 2 names with total storage of at least 4 kB. The four tiles of **Figure 1A** use 4 names, each naming a 1 kB tile, whereas the 2 pairs of tiles in **Figure 1B** can use 2 names to specify the paired tiles. In some embodiments, tile instructions accept a name of a paired tile as an operand. In operation, the tiles can be loaded from and stored to memory using load and store operations. Depending upon the instruction encoding scheme used, the amount of available application memory, as well as the size, number, and configuration of available tiles varies.

40 **[0016]** In some embodiments, tile parameters are definable. For example, a "palette" is used to provide tile options. Exemplary options include, but are not limited to: the number of tile names, the number of bytes in a row of storage, the number of rows and columns in a tile, etc. For example, a maximum "height" (number of rows) of a tile may be defined as:

45 **[0017]** $\text{Tile Max Rows} = \text{Architected Storage} / (\text{The Number of Palette Names} * \text{The Number of Bytes per row}).$

[0018] As such, an application can be written such that a fixed usage of names will be able to take advantage of different storage sizes across implementations.

50 **[0019]** Configuration of tiles is done using a matrix (tile) configuration ("TILECONFIG") instruction, where a particular

tile usage is defined in a selected palette. This declaration includes the number of tile names to be used, the requested number of rows and columns per name (tile), and, in some embodiments, the requested datatype of each tile. In some embodiments, consistency checks are performed during the execution of a TILECONFIG instruction to determine that it matches the restrictions of the palette entry.

EXEMPLARY TILE STORAGE TYPES

[0020] Figure 2 illustrates several examples of matrix storage. In (A), a tile is stored in memory. As shown, each "row" consists of four packed data elements. To get to the next "row," a stride value is used. Note that rows may be consecutively stored in memory. Strided memory accesses allow for access of one row to then next when the tile storage does not map the underlying memory array row width.

[0021] Tile loads from memory and stores to memory are typically strided accesses from the application memory to packed rows of data. Exemplary TILELOAD and TILES TORE instructions, or other instruction references to application memory as a TILE operand in load-op instructions, are, in some embodiments, restartable to handle (up to) 2*rows of page faults, unmasked floating-point exceptions, and/or interrupts per instruction.

[0022] In (B), a matrix is stored in a tile comprised of a plurality of registers such as packed data registers (single instruction, multiple data (SIMD) or vector registers). In this example, the tile is overlaid on three physical registers. Typically, consecutive registers are used, however, this need not be the case.

[0023] In (C), a matrix is stored in a tile in non-register storage accessible to a fused multiply accumulate (FMA) circuit used in tile operations. This storage may be inside of an FMA, or adjacent to it. Additionally, in some embodiments, discussed below, the storage may be for a data element and not an entire row or tile.

[0024] The supported parameters for the TMMA architecture are reported via CPUID. In some embodiments, the list of information includes a maximum height and a maximum STIVID dimension. Configuring the TMMA architecture requires specifying the dimensions for each tile, the element size for each tile and the palette identifier. This configuration is done by executing the TILECONFIG instruction.

[0025] Successful execution of a TILECONFIG instruction enables subsequent TILE operators. A TILERELASEALL instruction clears the tile configuration and disables the TILE operations (until the next TILECONFIG instructions executes). In some embodiments, XSAVE, XSTORE, etc. are used in context switching using tiles. In some embodiments, 2 XCR0 bits are used in XSAVE, one for TILECONFIG metadata and one bit corresponding to actual tile payload data.

[0026] TILECONFIG not only configures the tile usage, but also sets a state variable indicating that the program is in a region of code with tiles configured. An implementation may enumerate restrictions on other instructions that can be used with a tile region such as no usage of an existing register set, etc.

[0027] Exiting a tile region is typically done with the TILERELASEALL instruction. It takes no parameters and swiftly invalidates all tiles (indicating that the data no longer needs any saving or restoring) and clears the internal state corresponding to being in a tile region.

[0028] In some embodiments, tile operations will zero any rows and any columns beyond the dimensions specified by the tile configuration. For example, tile operations will zero the data beyond the configured number of columns (factoring in the size of the elements) as each row is written. For example, with 64-byte rows and a tile configured with 10 rows and 12 columns, an operation writing FP32 elements would write each of the first 10 rows with 12*4 bytes with output/result data and zero the remaining 4*4 bytes in each row. Tile operations also fully zero any rows after the first 10 configured rows. When using 1K tile with 64-byte rows, there would be 16 rows, so in this example, the last 6 rows would also be zeroed.

[0029] In some embodiments, a context restore instruction (e.g., XRSTOR), when loading data, enforces that the data beyond the configured rows for a tile will be maintained as zero. If there is no valid configuration, all rows are zeroed. XRSTOR of tile data can load garbage in the columns beyond those configured. It should not be possible for XRSTOR to clear beyond the number of columns configured because there is not an element width associated with the tile configuration.

[0030] Context save (e.g., XSAVE) exposes the entire TILE storage area when writing it to memory. If XRSTOR loaded garbage data into the rightmost part of a tile, that data will be saved by XSAVE. XSAVE will write zeros for rows beyond the number specified for each tile.

[0031] In some embodiments, tile instructions are restartable. The operations that access memory allow restart after page faults. The computational instructions that deal with floating-point operations also allow for unmasked floating-point exceptions, with the masking of the exceptions controlled by a control and/or status register.

[0032] To support restarting instructions after these events, the instructions store information in the start registers detailed below.

MATRIX (TILE) OPERATION SYSTEMS**EXEMPLARY HARDWARE SUPPORT**

5 **[0033]** **Figure 3** illustrates an embodiment of a system utilizing a matrix (tile) operations accelerator. In this illustration, a host processor/processing system 301 communicates commands 311 (e.g., matrix manipulation operations such as arithmetic or matrix manipulation operations, or load and store operations) to a matrix operations accelerator 307. However, this is shown this way for discussion purposes only. As detailed later, this accelerator 307 may be a part of a processing core. Typically, commands 311 that are tile manipulation operator instructions will refer to tiles as register-register ("reg-reg") or register-memory ("reg-mem") format. Other commands such as TILESTORE, TILELOAD, TILECONFIG, etc., do not perform data operations on a tile. Commands may be decoded instructions (e.g., micro-ops) or macro-instructions for the accelerator 307 to handle.

10 **[0034]** In this example, a coherent memory interface 303 is coupled to the host processor/processing system 301 and matrix operations accelerator 307 such that they can share memory. **Figures 4 and 5** show different embodiments of how memory is shared using a matrix operations accelerator. As shown in **Figure 4**, the host processor 401 and matrix operations accelerator circuitry 405 share the same memory 403. **Figure 5** illustrates an embodiment where the host processor 501 and matrix operations accelerator 505 do not share memory but can access each other's memory. For example, processor 501 can access tile memory 507 and utilize its host memory 503 as normal. Similarly, the matrix operations accelerator 505 can access host memory 503, but more typically uses its own memory 507. Note these memories may be of different types.

15 **[0035]** In some embodiments, tiles are supported using an overlay over physical registers. For example, a tile may utilize 16 1,024-bit registers, 32 512-bit registers, etc. depending on the implementation. In some embodiments, the matrix operations utilize 2-dimensional (2-D) data structures representing one or more packed regions of memory such as registers. Throughout this description, these 2-D data structures are referred to as tiles or tile registers.

20 **[0036]** In some embodiments, the matrix operations accelerator 307 includes a plurality of FMAs 309 coupled to data buffers 305 (in some implementations, one or more of these buffers 305 are stored in the FMAs of the grid as shown). The data buffers 305 buffer tiles loaded from memory and/or tiles to be stored to memory (e.g., using a tileload or tilestore instruction). Data buffers may be, for example, a plurality of registers. Typically, these FMAs are arranged as a grid of chained FMAs 309 which are able to read and write tiles. In this example, the matrix operations accelerator 307 is to perform a matrix multiply operation using tiles T0, T1, and T2. At least one of tiles is housed in the FMA grid 309. In some embodiments, all tiles in an operation are stored in the FMA grid 309. In other embodiments, only a subset is stored in the FMA grid 309. As shown, T1 is housed and T0 and T2 are not. Note that A, B, and C refer to the matrices of these tiles which may or may not take up the entire space of the tile.

25 **[0037]** **Figure 6** illustrates an embodiment of matrix multiply accumulate operation using tiles ("TMMA").

30 **[0038]** The number of rows in the matrix (TILE A 601) matches the number of serial (chained) FMAs comprising the computation's latency. An implementation is free to recirculate on a grid of smaller height, but the computation remains the same.

35 **[0039]** The source/destination vector comes from a tile of N rows (TILE C 605) and the grid of FMAs 611 performs N vector-matrix operations resulting in a complete instruction performing a matrix multiplication of tiles. Tile B 603 is the other vector source and supplies "broadcast" terms to the FMAs in each stage.

40 **[0040]** In operation, in some embodiments, the elements of matrix B (stored in a tile B 603) are spread across the rectangular grid of FMAs. Matrix B (stored in tile A 601) has its elements of a row transformed to match up with the columnar dimension of the rectangular grid of FMAs. At each FMA in the grid, an element of A and B are multiplied and added to the incoming summand (from above in the Figure) and the outgoing sum is passed to the next row of FMAs (or the final output).

45 **[0041]** The latency of a single step is proportional to K (row height of matrix B) and dependent TMMA's typically have enough source-destination rows (either in a single tile or across tile) to hide that latency. An implementation may also split the SIMD (packed data element) dimension M (row height of matrix A) across time steps, but this simply changes the constant that K is multiplied by. When a program specifies a smaller K than the maximum enumerated by the TMACC, an implementation is free to implement this with "masking" or "early outs."

50 **[0042]** The latency of an entire TMMA is proportional to $N \cdot K$. The repeat rate is proportional to N. The number of MACs per TMMA instruction is $N \cdot K \cdot M$.

55 **[0043]** **Figure 7** illustrates an embodiment of a subset of the execution of an iteration of a chained fused multiply accumulate instruction. In particular, this illustrates execution circuitry of an iteration of one packed data element position of the destination. In this embodiment, the chained fused multiply accumulate is operating on signed sources wherein the accumulator is 2x the input data size.

[0044] A first signed source (source 1 701) and a second signed source (source 2 703) each have four packed data elements. Each of these packed data elements stores signed data such as floating-point data. A third signed source

(source 3 709) has two packed data elements, each of which stores signed data. The sizes of the first and second signed sources 701 and 703 are half that of the third signed source (initial value or previous result) 709. For example, the first and second signed sources 701 and 703 could have 32-bit packed data elements (e.g., single precision floating-point) while the third signed source 709 could have 64-bit packed data elements (e.g., double precision floating-point).

5 **[00445]** In this illustration, only the two most significant packed data element positions of the first and second signed sources 701 and 703 and the most significant packed data element position of the third signed source 709 are shown. Of course, the other packed data element positions would also be processed.

10 **[00446]** As illustrated, packed data elements are processed in pairs. For example, the data of the most significant packed data element positions of the first and second signed sources 701 and 703 are multiplied using a multiplier circuit 705, and the data from second most significant packed data element positions of the first and second signed sources 701 and 703 are multiplied using a multiplier circuit 707. In some embodiments, these multiplier circuits 705 and 707 are reused for other packed data elements positions. In other embodiments, additional multiplier circuits are used so that the packed data elements are processed in parallel. In some contexts, parallel execution is done using lanes that are the size of the signed third source 709. The results of each of the multiplications are added using addition circuitry 711.

15 **[00447]** The result of the addition of the results of the multiplications is added to the data from most significant packed data element position of the signed source 3 709 (using a different adder 713 or the same adder 711).

20 **[00448]** Finally, the result of the second addition is either stored into the signed destination 715 in a packed data element position that corresponds to the packed data element position used from the signed third source 709 or passed on to the next iteration if there is one. In some embodiments, a writemask is applied to this storage such that if a corresponding writemask (bit) is set, the storage happens, and, if not set, the storage does not happen.

[00449] Figure 8 illustrates an embodiment of a subset of the execution of an iteration of a chained fused multiply accumulate instruction. In particular, this illustrates execution circuitry of an iteration of one packed data element position of the destination. In this embodiment, the chained fused multiply accumulate is operating on signed sources wherein the accumulator is 2x the input data size.

25 **[00500]** A first signed source (source 1 801) and a second signed source (source 2 803) each have four packed data elements. Each of these packed data elements stores signed data such as integer data. A third signed source (source 3 809) has two packed data elements, each of which stores signed data. The sizes of the first and second signed sources 801 and 803 are half that of the third signed source 809. For example, the first and second signed sources 801 and 803 could have 32-bit packed data elements (e.g., single precision floating-point) the third signed source 809 could have

30 64-bit packed data elements (e.g., double precision floating-point). **[00501]** In this illustration, only the two most significant packed data element positions of the first and second signed sources 801 and 803 and the most significant packed data element position of the third signed source 809 are shown. Of course, the other packed data element positions would also be processed.

35 **[00502]** As illustrated, packed data elements are processed in pairs. For example, the data of the most significant packed data element positions of the first and second signed sources 801 and 803 are multiplied using a multiplier circuit 805, and the data from second most significant packed data element positions of the first and second signed sources 801 and 803 are multiplied using a multiplier circuit 807. In some embodiments, multiplier circuits 805 and 807 perform the multiplications with infinite precision without saturation and use adder/saturation circuitry 813 to saturate the results of the accumulation to plus or minus infinity in case of an overflow and to zero in case of any underflow. In other

40 **[00503]** Addition/saturation (accumulator) circuitry 813 preserves a sign of an operand when the addition results in a value that is too big. In particular, saturation evaluation occurs on the infinite precision result between the multi-way-add and the write to the destination or next iteration. When the accumulator 813 is floating-point and the input terms are integer, the sum of products and the floating-point accumulator input value are turned into infinite precision values (fixed point numbers of hundreds of bits), the addition of the multiplication results and the third input is performed, and a single

50 rounding to the actual accumulator type is performed. **[00504]** Unsigned saturation means the output values are limited to a maximum unsigned number for that element width (all 1s). Signed saturation means a value is limited to be in the range between a minimum negative number and a max positive number for that element width (for bytes for example, the range is from $-128 (= -2^7)$ to $127 (= 2^7 - 1)$).

55 **[00505]** The result of the addition and saturation check is stored into the signed result 815 in a packed data element position that corresponds to the packed data element position used from the signed third source 809 or passed on to the next iteration if there is one. In some embodiments, a writemask is applied to this storage such that if a corresponding writemask (bit) is set, the storage happens, and, if not set, the storage does not happen.

[00506] Figure 9 illustrates an embodiment of a subset of the execution of an iteration of a chained fused multiply

accumulate instruction. In particular, this illustrates execution circuitry of an iteration of one packed data element position of the destination. In this embodiment, the chained fused multiply accumulate is operating on a signed source and an unsigned source wherein the accumulator is 4x the input data size.

5 [0057] A first signed source (source 1 901) and a second unsigned source (source 2 903) each have four packed data elements. Each of these packed data elements has data such as floating-point or integer data. A third signed source (initial value or result 915) has a packed data element of which stores signed data. The sizes of the first and second sources 901 and 903 are a quarter of the third signed source 915. For example, the first and second sources 901 and 903 could have 16-bit packed data elements (e.g., word) and the third signed source 915 could have 64-bit packed data elements (e.g., double precision floating-point or 64-bit integer).

10 [0058] In this illustration, the four most significant packed data element positions of the first and second sources 901 and 903 and the most significant packed data element position of the third signed source 915 are shown. Of course, other packed data element positions would also be processed if there are any.

15 [0059] As illustrated, packed data elements are processed in quadruplets. For example, the data of the most significant packed data element positions of the first and second sources 901 and 903 are multiplied using a multiplier circuit 905, data from second most significant packed data element positions of the first and second sources 901 and 903 are multiplied using a multiplier circuit 907, data from third most significant packed data element positions of the first and second sources 901 and 903 are multiplied using a multiplier circuit 909, and data from the least significant packed data element positions of the first and second sources 901 and 903 are multiplied using a multiplier circuit 911. In some embodiments, the signed packed data elements of the first source 901 are sign extended and the unsigned packed data elements of the second source 903 are zero extended prior to the multiplications.

20 [0060] In some embodiments, these multiplier circuits 905-911 are reused for other packed data elements positions. In other embodiments, additional multiplier circuits are used so that the packed data elements are processed in parallel. In some contexts, parallel execution is done using lanes that are the size of the signed third source 915. The results of each of the multiplications are added using addition circuitry 913.

25 [0061] The result of the addition of the results of the multiplications is added to the data from most significant packed data element position of the signed source 3 915 (using a different adder 917 or the same adder 913).

30 [0062] Finally, the result 919 of the second addition is either stored into the signed destination in a packed data element position that corresponds to the packed data element position used from the signed third source 915 or passed to the next iteration. In some embodiments, a writemask is applied to this storage such that if a corresponding writemask (bit) is set, the storage happens, and, if not set, the storage does not happen.

[0063] **Figure 10** illustrates an embodiment of a subset of the execution of an iteration of chained fused multiply accumulate instruction. In particular, this illustrates execution circuitry of an iteration of one packed data element position of the destination. In this embodiment, the chained fused multiply accumulate is operating on a signed source and an unsigned source wherein the accumulator is 4x the input data size.

35 [0064] A first signed source 1001 and a second unsigned source 1003 each have four packed data elements. Each of these packed data elements stores data such as floating-point or integer data. A third signed source 1015 (initial or previous result) has a packed data element of which stores signed data. The sizes of the first and second sources are a quarter of the third signed source 1015 (initial or previous result). For example, the first and second sources could have 16-bit packed data elements (e.g., word) and the third signed source 1015 (initial or previous result) could have 64-bit packed data elements (e.g., double precision floating-point or 64-bit integer).

40 [0065] In this illustration, the four most significant packed data element positions of the first signed source 1001 and the second unsigned source 1003 and the most significant packed data element position of the third signed source 1015 are shown. Of course, other packed data element positions would also be processed if there are any.

45 [0066] As illustrated, packed data elements are processed in quadruplets. For example, the data of the most significant packed data element positions of the first signed source 1001 and the second unsigned source 1003 are multiplied using a multiplier circuit 1005, data from second most significant packed data element positions of the first signed source 1001 and the second unsigned source 1003 are multiplied using a multiplier circuit 1007, data from third most significant packed data element positions of the first signed source 1001 and the second unsigned source 1003 are multiplied using a multiplier circuit 1009, and data from the least significant packed data element positions of the first signed source 1001 and the second unsigned source 1003 are multiplied using a multiplier circuit 1011. In some embodiments, the signed packed data elements of the first signed source 1001 are sign extended and the unsigned packed data elements of the second unsigned source 1003 are zero extended prior to the multiplications.

50 [0067] In some embodiments, these multiplier circuits 1005-1011 are reused for other packed data elements positions. In other embodiments, additional multiplier circuits are used so that the packed data elements are processed in parallel. In some contexts, parallel execution is done using lanes that are the size of third signed source 1015 (initial or previous result). The result of the addition of the results of the multiplications is added to the data from most significant packed data element position of third signed source 1015 (initial or previous result) using adder/saturation 1013 circuitry.

55 [0068] Addition/saturation (accumulator) circuitry 1013 preserves a sign of an operand when the addition results in a

value that is too big or too small for signed saturation. In particular, saturation evaluation occurs on the infinite precision result between the multi-way-add and the write to the destination. When the accumulator 1013 is floating-point and the input terms are integer, the sum of products and the floating-point accumulator input value are turned into infinite precision values (fixed point numbers of hundreds of bits), the addition of the multiplication results and the third input is performed, and a single rounding to the actual accumulator type is performed.

[0069] The result 1019 of the addition and saturation check is stored into the signed destination in a packed data element position that corresponds to the packed data element position used from third signed source 1015 (initial or previous result) or passed to the next iteration. In some embodiments, a writemask is applied to this storage such that if a corresponding writemask (bit) is set, the storage happens, and, if not set, the storage does not happen.

[0070] **Figure 11** illustrates power-of-two sized SIIVID implementations wherein the accumulators use input sizes that are larger than the inputs to the multipliers according to an embodiment. Note the source (to the multipliers) and accumulator values may be signed or unsigned values. For an accumulator having 2X input sizes (in other words, the accumulator input value is twice the size of the packed data element sizes of the sources), table 1101 illustrates different configurations. For byte sized sources, the accumulator uses word or half-precision floating-point (HPFP) values that are 16-bit in size. For word sized sources, the accumulator uses 32-bit integer or single-precision floating-point (SPFP) values that are 32-bit in size. For SPFP or 32-bit integer sized sources, the accumulator uses 64-bit integer or double-precision floating-point (DPFP) values that are 64-bit in size.

[0071] For an accumulator having 4X input sizes (in other words, the accumulator input value is four times the size of the packed data element sizes of the sources), table 1103 illustrates different configurations. For byte sized sources, the accumulator uses 32-bit integer or single-precision floating-point (SPFP) values that are 32-bit in size. For word sized sources, the accumulator uses 64-bit integer or double-precision floating-point (DPFP) values that are 64-bit in size in some embodiments.

[0072] For an accumulator having 8X input sizes (in other words, the accumulator input value is eight times the size of the packed data element sizes of the sources), table 1105 illustrates a configuration. For byte sized sources, the accumulator uses 64-bit integer.

[0073] As hinted at earlier, matrix operations circuitry may be included in a core, or as an external accelerator. **Figure 12** illustrates an embodiment of a system utilizing matrix operations circuitry. In this illustration, multiple entities are coupled with a ring interconnect 1245.

[0074] A plurality of cores, core 0 1201, core 1 1203, core 2 1205, and core N 1207 provide non-tile-based instruction support. In some embodiments, matrix operations circuitry 1251 is provided in a core 1203, and in other embodiments matrix operations circuitry 1211 and 1213 are accessible on the ring interconnect 1245.

[0075] Additionally, one or more memory controllers 1223-1225 are provided to communicate with memory 1233 and 1231 on behalf of the cores and/or matrix operations circuitry.

[0076] **Figure 13** illustrates an embodiment of a processor core pipeline supporting matrix operations using tiles. Branch prediction and decode circuitry 1303 performs branch predicting of instructions, decoding of instructions, and/or both from instructions stored in instruction storage 1301. For example, instructions detailed herein may be stored in instruction storage. In some implementations, separate circuitry is used for branch prediction and in some embodiments, at least some instructions are decoded into one or more micro-operations, microcode entry points, microinstructions, other instructions, or other control signals using microcode 1305. The branch prediction and decode circuitry 1303 may be implemented using various different mechanisms. Examples of suitable mechanisms include, but are not limited to, look-up tables, hardware implementations, programmable logic arrays (PLAs), microcode read only memories (ROMs), etc.

[0077] The branch prediction and decode circuitry 1303 is coupled to allocate/rename 1307 circuitry which is coupled, in some embodiments, to scheduler circuitry 1309. In some embodiments, these circuits provide register renaming, register allocation, and/or scheduling functionality by performing one or more of: 1) renaming logical operand values to physical operand values (e.g., a register alias table in some embodiments), 2) allocating status bits and flags to the decoded instruction, and 3) scheduling the decoded instruction for execution on execution circuitry out of an instruction pool (e.g., using a reservation station in some embodiments).

[0078] The scheduler circuitry 1309 represents any number of different schedulers, including reservations stations, central instruction window, etc. The scheduler circuitry 1309 is coupled to, or includes, physical register file(s) 1315. Each of the physical register file(s) 1315 represents one or more physical register files, different ones of which store one or more different data types, such as scalar integer, scalar floating-point, packed integer, packed floating-point, vector integer, vector floating-point, status (e.g., an instruction pointer that is the address of the next instruction to be executed), tiles, etc. In one embodiment, the physical register file(s) 1315 comprises vector registers circuitry, write mask registers circuitry, and scalar registers circuitry. These register circuits may provide architectural vector registers, vector mask registers, and general-purpose registers. The physical register file(s) 1315 is overlapped by a retirement circuit 1317 to illustrate various ways in which register renaming and out-of-order execution may be implemented (e.g., using a reorder buffer(s) and a retirement register file(s); using a future file(s), a history buffer(s), and a retirement register

file(s); using a register maps and a pool of registers; etc.). The retirement circuit 1317 and the physical register file(s) 1315 are coupled to the execution circuitry 1311.

[0079] While register renaming is described in the context of out-of-order execution, it should be understood that register renaming may be used in an in-order architecture. While the illustrated embodiment of the processor may also include separate instruction and data cache units and a shared L2 cache unit, alternative embodiments may have a single internal cache for both instructions and data, such as, for example, a Level 1 (L1) internal cache, or multiple levels of internal cache. In some embodiments, the system may include a combination of an internal cache and an external cache that is external to the core and/or the processor. Alternatively, all of the cache may be external to the core and/or the processor.

[0080] The execution circuitry 1311 is a set of one or more execution circuits, including scalar circuitry 1321, vector/SIMD circuitry 1323, and matrix operations circuitry 1327, as well as memory access circuitry 1325 to access cache 1313. The execution circuits perform various operations (e.g., shifts, addition, subtraction, multiplication) and on various types of data (e.g., scalar floating-point, packed integer, packed floating-point, vector integer, vector floating-point). While some embodiments may include a number of execution units dedicated to specific functions or sets of functions, other embodiments may include only one execution unit or multiple execution units that all perform all functions. The scalar circuitry 1321 performs scalar operations, the vector/SIMD circuitry 1323 performs vector/SIMD operations, and matrix operations circuitry 1327 performs matrix (tile) operations detailed herein.

[0081] By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement a pipeline as follows: 1) an instruction fetch circuit performs fetch and length decoding stages; 2) the branch and decode circuitry 1303 performs a decode stage; 3) the allocate/rename 1307 circuitry performs an allocation stage and renaming stage; 4) the scheduler circuitry 1309 performs a schedule stage; 5) physical register file(s) (coupled to, or included in, the scheduler circuitry 1309 and allocate/rename 1307 circuitry and a memory unit perform a register read/memory read stage; the execution circuitry 1311 performs an execute stage; 6) a memory unit and the physical register file(s) unit(s) perform a write back/memory write stage; 7) various units may be involved in the exception handling stage; and 8) a retirement unit and the physical register file(s) unit(s) perform a commit stage.

[0082] The core may support one or more instructions sets (e.g., the x86 instruction set (with some extensions that have been added with newer versions); the MIPS instruction set of MIPS Technologies of Sunnyvale, CA; the ARM instruction set (with optional additional extensions such as NEON) of ARM Holdings of Sunnyvale, CA), including the instruction(s) described herein. In one embodiment, the core 1390 includes logic to support a packed data instruction set extension (e.g., AVX1, AVX2), thereby allowing the operations used by many multimedia applications to be performed using packed data.

[0083] It should be understood that the core may support multithreading (executing two or more parallel sets of operations or threads), and may do so in a variety of ways including time sliced multithreading, simultaneous multithreading (where a single physical core provides a logical core for each of the threads that physical core is simultaneously multithreading), or a combination thereof (e.g., time sliced fetching and decoding and simultaneous multithreading thereafter such as in the Intel® Hyperthreading technology).

[0084] Figure 14 illustrates an embodiment of a processor core pipeline supporting matrix operations using tiles. Branch prediction and decode circuitry 1403 performs branch predicting of instructions, decoding of instructions, and/or both from instructions stored in instruction storage 1401. For example, instructions detailed herein may be stored in instruction storage. In some implementations, separate circuitry is used for branch prediction and in some embodiments, at least some instructions are decoded into one or more micro-operations, microcode entry points, microinstructions, other instructions, or other control signals using microcode 1405. The branch prediction and decode circuitry 1403 may be implemented using various different mechanisms. Examples of suitable mechanisms include, but are not limited to, look-up tables, hardware implementations, programmable logic arrays (PLAs), microcode read only memories (ROMs), etc.

[0085] The branch prediction and decode circuitry 1403 is coupled to allocate/rename 1407 circuitry which is coupled, in some embodiments, to scheduler circuitry 1409. In some embodiments, these circuits provide register renaming, register allocation, and/or scheduling functionality by performing one or more of: 1) renaming logical operand values to physical operand values (e.g., a register alias table in some embodiments), 2) allocating status bits and flags to the decoded instruction, and 3) scheduling the decoded instruction for execution on execution circuitry out of an instruction pool (e.g., using a reservation station in some embodiments).

[0086] The scheduler circuitry 1409 represents any number of different schedulers, including reservations stations, central instruction window, etc. The scheduler unit(s) scheduler circuitry 1409 is coupled to, or includes, physical register file(s) 1415. Each of the physical register file(s) 1415 represents one or more physical register files, different ones of which store one or more different data types, such as scalar integer, scalar floating-point, packed integer, packed floating-point, vector integer, vector floating-point, status (e.g., an instruction pointer that is the address of the next instruction to be executed), tiles, etc. In one embodiment, the physical register file(s) 1415 comprises vector registers circuitry, write mask registers circuitry, and scalar registers circuitry. These register circuits may provide architectural vector

registers, vector mask registers, and general-purpose registers. The physical register file(s) 1415 is overlapped by a retirement circuit 1417 to illustrate various ways in which register renaming and out-of-order execution may be implemented (e.g., using a reorder buffer(s) and a retirement register file(s); using a future file(s), a history buffer(s), and a retirement register file(s); using a register maps and a pool of registers; etc.). The retirement circuit 1417 and the physical register file(s) 1415 are coupled to the execution circuitry 1411.

[0087] While register renaming is described in the context of out-of-order execution, it should be understood that register renaming may be used in an in-order architecture. While the illustrated embodiment of the processor may also include separate instruction and data cache units and a shared L2 cache unit, alternative embodiments may have a single internal cache for both instructions and data, such as, for example, a Level 1 (L1) internal cache, or multiple levels of internal cache. In some embodiments, the system may include a combination of an internal cache and an external cache that is external to the core and/or the processor. Alternatively, all of the cache may be external to the core and/or the processor.

[0088] The execution circuitry 1411 a set of one or more execution circuits 1427 and a set of one or more memory access circuits 1425 to access cache 1413. The execution circuits 1427 perform matrix (tile) operations detailed herein.

[0089] By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement a pipeline as follows: 1) an instruction fetch circuit performs fetch and length decoding stages; 2) the branch and decode circuitry 1403 performs a decode stage; 3) the allocate/rename 1407 circuitry performs an allocation stage and renaming stage; 4) the scheduler circuitry 1409 performs a schedule stage; 5) physical register file(s) (coupled to, or included in, the scheduler circuitry 1409 and allocate/rename 1407 circuitry and a memory unit perform a register read/memory read stage; the execution circuitry 1411 performs an execute stage; 6) a memory unit and the physical register file(s) unit(s) perform a write back/memory write stage; 7) various units may be involved in the exception handling stage; and 8) a retirement unit and the physical register file(s) unit(s) perform a commit stage.

[0090] The core may support one or more instructions sets (e.g., the x86 instruction set (with some extensions that have been added with newer versions); the MIPS instruction set of MIPS Technologies of Sunnyvale, CA; the ARM instruction set (with optional additional extensions such as NEON) of ARM Holdings of Sunnyvale, CA), including the instruction(s) described herein. In one embodiment, the core 1490 includes logic to support a packed data instruction set extension (e.g., AVX1, AVX2), thereby allowing the operations used by many multimedia applications to be performed using packed data.

[0091] It should be understood that the core may support multithreading (executing two or more parallel sets of operations or threads), and may do so in a variety of ways including time sliced multithreading, simultaneous multithreading (where a single physical core provides a logical core for each of the threads that physical core is simultaneously multithreading), or a combination thereof (e.g., time sliced fetching and decoding and simultaneous multithreading thereafter such as in the Intel® Hyperthreading technology).

LAYOUT

[0092] Throughout this description, data is expressed using row major data layout. Column major users should translate the terms according to their orientation. **Figure 15** illustrates an example of a matrix expressed in row major format and column major format. As shown, matrix A is a 2x3 matrix. When this matrix is stored in row major format, the data elements of a row are consecutive. When this matrix is stored in column major format, the data elements of a column are consecutive. It is a well-known property of matrices that $A^T * B^T = (BA)^T$, where superscript T means transform. Reading column major data as row major data results in the matrix looking like the transform matrix.

[0093] In some embodiments, row-major semantics are utilized in hardware, and column major data is to swap the operand order with the result being transforms of matrix, but for subsequent column-major reads from memory it is the correct, non-transformed matrix.

[0094] For example, if there are two column-major matrices to multiply:

$$\begin{array}{ccc}
 \begin{matrix} a & b \\ c & d \\ e & f \end{matrix} & * & \begin{matrix} g & i & k \\ h & j & l \end{matrix} \\
 (3 \times 2) & & (2 \times 3)
 \end{array} = \begin{array}{ccc}
 ag+bh & ai+bj & ak+bl \\
 cg+dh & ci+dj & ck+dl \\
 eg+fh & ei+fj & ek+fl \\
 (3 \times 3)
 \end{array}$$

[0095] The input matrices would be stored in linear memory (column-major) as:

EP 4 276 608 A2

a c e b d f

and

g h i j k l.

[0096] Reading those matrices as row-major with dimensions 2x3 and 3x2, they would appear as:

a c e and g h
b d f i j
k l

[0097] Swapping the order and matrix multiplying:

g h a c e ag+bh cg+dh eg+fh
i j * b d f = ai+bj ci+dj ei+fj
k l ak+bl ck+dl ek+fl

[0098] The transform matrix is out and can then be stored in in row-major order:

ag+bh cg+dh eg+fh ai+bj ci+dj ei+fj ak+bl ck+dl ek+fl

[0099] and used in subsequent column major computations, it is the correct un-transformed matrix:

ag+bh ai+bj ak+bl
cg+dh ci+dj ck+dl
eg+fh ei+fj ek+fl

EXEMPLARY USAGE

[0100] **Figure 16** illustrates an example of usage of matrices (e.g., tiles). In this example, matrix C 1601 includes two tiles, matrix A 1603 includes one tile, and matrix B 1605 includes two tiles. This figure shows an example of the inner loop of an algorithm to compute a matrix multiplication. In this example, two result tiles, tmm0 and tmm1, from matrix C 1601 are used to accumulate the intermediate results. One tile from the matrix A 1603 (tmm2) is re-used twice as it multiplied by two tiles from matrix B 1605. Pointers to load a new A matrix (tile) and two new B matrices (e.g., tiles) from the directions indicated by the arrows. An outer loop, not shown, adjusts the pointers for the C tiles.

[0101] The exemplary code as shown includes the usage of a tile configuration instruction and is executed to configure tile usage, load tiles, a loop to process the tiles, store tiles to memory, and release tile usage.

[0102] **Figure 17** illustrates an embodiment of usage of matrices (e.g., tiles). At 1701, tile usage is configured. For example, a TILECONFIG instruction is executed to configure tile usage including setting a number of rows and columns per tile. Typically, at least one matrix (tile) is loaded from memory at 1703. At least one matrix (tile) operation is performed at 1705 using the matrices (e.g., tiles). At 1707, at least one matrix (tile) is stored out to memory and a context switch can occur at 1709.

EXEMPLARY CONFIGURATION

TILE CONFIGURATION HARDWARE SUPPORT

[0103] As discussed above, tile usage typically needs to be configured prior to use. For example, full usage of all rows and columns may not be needed. Not only does not configuring these rows and columns save power in some embodiments, but the configuration may be used to determine if an operation will generate an error. For example, a matrix

multiplication of the form $(N \times M) * (L \times N)$ will typically not work if M and L are not the same.

[0104] Prior to using matrices using tiles, in some embodiments, tile support is to be configured. For example, how many rows and columns per tile, tiles that are to be used, etc. are configured. A TILECONFIG instruction is an improvement to a computer itself as it provides for support to configure the computer to use a matrix accelerator (either as a part of a processor core, or as an external device). In particular, an execution of the TILECONFIG instruction causes a configuration to be retrieved from memory and applied to matrix (tile) settings within a matrix accelerator.

TILE USAGE CONFIGURATION

[0105] Figure 18 illustrates support for configuration of the usage of tiles according to an embodiment. A memory 1801 contains the tile description 1803 of the matrices (e.g., tiles) to be supported.

[0106] Instruction execution resources 1811 of a processor/core 1805 stores aspects of a tile description 1803 into tile configurations 1817. The tile configurations 1817 include palette table 1813 to detail what tiles for a palette are configured (the number of rows and columns in each tile) and a marking that matrix support is in use. In particular, instruction execution resources 1811 are configured to use tiles as specified by the tile configurations 1817. The instruction execution resources 1811 may also include a machine specific register or configuration register to indicate tile usage. Additional values such as in-use and start values are also set. The tile configurations 1817 utilize register(s) 1819 to store tile usage and configuration information.

[0107] Figure 19 illustrates an embodiment of a description of the matrices (e.g., tiles) to be supported. This is the description that is to be stored upon an execution of a STTILECFG instruction. In this example, each field is a byte. In byte [0], a palette ID 1901 is stored. The palette ID is used to index a palette table 1813 which stores, per palette ID, a number of bytes in a tile, and bytes per row of the tiles that are associated with this ID as defined by the configuration.

[0108] Byte 1 stores a value to be stored in a "startRow" register 1903 and byte 2 stores a value to be stored in a register, startP 1905. To support restarting instructions after these events, the instructions store information these registers. To support restarting instructions after break events such as those detailed above, the instructions store information in these registers. The startRow value indicates the row that should be used for restart. The startP value indicates the position within the row for store operations when pairs are used and, in some embodiments, indicates the lower half of the row (in the lower tile of a pair) or higher half of the row (in the higher tile of a pair). Generally, the position in the row (the column) is not needed.

[0109] With the exception of TILECONFIG and STTILECFG, successfully executing matrix (tile) instructions will set both startRow and startP to zero.

[0110] Any time an interrupted matrix (tile) instruction is not restarted, it is the responsibility of software to zero the startRow and startP values. For example, unmasked floating-point exception handlers might decide to finish the operation in software and change the program counter value to another instruction, usually the next instruction. In this case the software exception handler must zero the startRow and startP values in the exception presented to it by the operating system before resuming the program. The operating system will subsequently reload those values using a restore instruction.

[0111] Byte 3 stores an indication of pairs (1b per tile) of tiles 1907.

[0112] Bytes 16-17 store the number of rows 1913 and columns 1915 for tile 0, bytes 18-19 store the number of rows and columns for tile 1, etc. In other words, each 2-byte group specifies a number of rows and columns for a tile. If a group of 2 bytes is not used to specify tile parameters, they should have the value zero. Specifying tile parameters for more tiles than the implementation limit or the palette limit results in a fault. Unconfigured tiles are set to an initial state with 0 rows, 0 columns.

[0113] Finally, the configuration in memory typically ends with an ending delineation such as all zeros for several consecutive bytes.

EXEMPLARY TILE AND TILE CONFIGURATION STORAGE

[0114] Figures 20(A)-(D) illustrate examples of register(s) 1819. Figure 20(A) illustrates a plurality of registers 1819. As shown each tile (TMM0 2001 ... TMMN 2003) has a separate register with each register storing a row and column size for that particular tile. StartP 2011 and StartRow 2013 are stored in separate registers. One or more status registers 2015 are set (e.g., TILES_CONFIGURED = 1) to indicate tiles are configured for use.

[0115] Figure 20(B) illustrates a plurality of registers 1819. As shown each tile has separate registers for its rows and columns. For example, TMM0 rows configuration 2021, TMM0 columns configuration 2023, StartP 2011 and StartRow 2013 are stored in separate registers. One or more status registers 2015 are set (e.g., TILES_CONFIGURED = 1) to indicate tiles are configured for use.

[0116] Figure 20(C) illustrates a single register 1819. As shown, this register stores tile configurations (rows and columns per tile) 2031, StartP 2011, and StartRow 2013 are stored in single register as packed data registers. One or

more status registers 2015 are set (e.g., TILES_CONFIGURED = 1) to indicate tiles are configured for use.

[0117] Figure 20(D) illustrates a plurality of registers 1819. As shown, a single register stores tile configuration (rows and columns per tile) 2031. StartP and StartRow are stored in separate registers 2011 and 2013. One or more status registers 2015 are set (e.g., TILES_CONFIGURED = 1) to indicate tiles are configured for use.

[0118] Other combinations are contemplated such as combining the start registers into a single register where they are shown separately, etc.

TDPBF8PS

[0119] As mentioned above, special hardware for General Matrix Multiplication (a.k.a., GEMM), is a good option for improving the peak compute performance (and energy efficiency) of certain applications, such as deep learning. Some of these applications, including deep learning, can operate on input data elements with relatively few bits without losing accuracy, as long as the output elements have enough bits (i.e., more than the inputs).

[0120] Accordingly, disclosed methods and systems perform an 8-bit floating-point matrix dot product operation TILED PBF8PS (TDPBF8PS) that takes source matrices (e.g., tiles) having 8-bit floating-point elements, performs dot product multiplications, and accumulates the resulting products with a 32-bit single-precision destination.

[0121] In certain embodiments, an 8-bit floating point format is an eight-bit wide brain floating point format (BF8), which is an Institute of Electrical and Electronics Engineers (IEEE) (e.g., IEEE 754 standard) half-precision binary floating-point format (IEEE FP16) with the lower half (8 LSBs) cut off. The BF8 format may include a sign field (one bit wide), an exponent field (five bits wide), and a mantissa (significand precision) field (two bit wide). In certain embodiments, a mantissa (significand precision) field is presumed to have an implicit leading bit with value of one, unless the exponent field is stored with all zeros. Further, a 32-bit floating-point format may include binary32 (according to an IEEE standard), which is sometimes referred to herein as "single-precision" or "fp32", e.g., having a sign field (one bit wide), an exponent field (eight bits wide), and a mantissa (significand precision) field (twenty four bits implicitly stored, i.e., twenty three bits wide explicitly stored).

[0122] In certain embodiments, the disclosed TDPBF8PS instruction is to be executed by a processor that includes fetch circuitry to fetch an instruction having fields to specify an opcode and locations of a M by N destination matrix (tile) having single-precision elements, a M by K first source matrix (tile), and a K by N second source matrix (tile), each element of the specified first and second source matrices including a quadruple of 8-bit floating point values including first (e.g., having a quad index of zero in the (0,1,2,3) quads in each element of the sources in Figure 21A), second (e.g., having a quad index of one in the (0,1,2,3) quads in each element of the sources in Figure 21A), third (e.g., having a quad index of two in the (0,1,2,3) quads in each element of the sources in Figure 21A), and fourth (e.g., having a quad index of three in the (0,1,2,3) quads in each element of the sources in Figure 21A) 8-bit floating-point values, wherein the opcode is to indicate execution circuitry is to, for each element (e.g., each of MxN number of elements) of the specified destination matrix (e.g., tile), convert K quadruples of 8-bit floating point values from row M of the specified first source matrix (e.g., tile) and K corresponding quadruples of 8-bit floating point values from column N of the specified second source matrix (e.g., tile) to single-precision values, multiply K pairs of the converted first values from the two specified source matrices (e.g., tiles) to generate K first products, multiply K pairs of the converted second values from the two specified source matrices (e.g., tiles) to generate K second products, multiply K pairs of the converted third values from the two specified source matrices (e.g., tiles) to generate K third products, multiply K pairs the converted fourth values from the two specified source matrices (e.g., tiles) to generate K fourth products, accumulate the first and second products to generate a first accumulated sum, separately accumulate the third and fourth products to generate a second accumulated sum, add the first and second accumulated sums to a final accumulated sum to be added with the previous contents of element (M,N).

[0123] In certain embodiments, the disclosed TDPBF8PS instruction is to be executed by a processor that includes fetch circuitry to fetch an instruction having fields to specify an opcode that indicates execution circuitry is to cause, for each element of the first source matrix and corresponding element of the second source matrix, a conversion of the 8-bit floating-point values to single-precision values, a multiplication of converted single-precision values from first values of the quads together to generate a first result, a multiplication of converted single-precision values from second values of the quads together to generate a second result, a multiplication of converted single-precision values from third values of the quads together to generate a third result, a multiplication of converted single-precision values from fourth values of the quads together to generate a fourth result, and an accumulation of the first, second, third, and fourth results with previous contents of a corresponding element of the destination matrix.

[0124] In certain embodiments, the processor will also include other supporting hardware, such as decode circuitry to decode the fetched instruction, and execution circuitry to respond to the decoded instruction as specified by the opcode, e.g., execution circuitry that causes a matrix operations accelerator (e.g., matrix operations accelerator 307 in Figure 3) to perform one or more (e.g., all) of the actions of the TDPBF8PS instruction.

[0125] Figure 21A is a block diagram illustrating use of a TDPBF8PS instruction to accelerate matrix multiplication,

according to some embodiments. As shown, instruction 2101 includes fields to specify an opcode 2102 (e.g., TDPBF8PS) and locations of a M by N destination matrix (e.g., tile) 2104 having single-precision elements, a M by K first source matrix (e.g., tile) 2106, and a K by N second source matrix (e.g., tile) 2108, the specified source matrices having elements that each comprise a quadruple of 8-bit (e.g., BF8) floating-point values. A format of the TDPBF8PS instruction, according

5 to some embodiments, is further illustrated and described at least with respect to **Figures 24, 25A-B, and 26A-D.**
[0126] Here, the specified first source matrix (e.g., tile) 2112A has dimensions of M=4 by K=3. The specified second source matrix (e.g., tile) 2112B has dimensions of K=3 by N=5. K, M, and N are shown as having different values for illustrative purposes, but in other embodiments, they can be equal.

10 **[0127]** In one embodiment of operation, processor 2100 is to respond to opcode 2102 (TDPBF8PS) by, for each element (M,N) of the specified destination matrix (e.g., tile) 2122, convert, using convert circuit 2116A, K quads of values from row M of the specified first source matrix (e.g., tile) 2112A, and, using convert circuit 2116B, K quads of values from column N of the specified second source matrix (e.g., tile) 2112B to single-precision, e.g., binary32 single-precision floating-point as specified by IEEE 794. Processor 2100 is then to multiply, using multiply circuit 2118, the K pairs of converted first quad values together, the K pairs of converted second quad values together, the K pairs of converted third quad values together, and the K pairs of converted fourth quad values together, and accumulate, using accumulate circuit 2120, the products (4*K number of products) with previous contents of the element (M,N).

15 **[0128]** Performance of the TDPBF8PS instruction is illustrated here for setting the destination element at matrix (e.g., tile) location, e.g., the index of (row, column) being (1, 0) in $C_{1,0}$. In Figure 21A, the ".0" refers to the first value of a quad, the ".1" refers to the second value of the quad, the ".2" refers to the third value of the quad, the ".3" refers to the fourth value of the quad, e.g., such that $A_{1,0,0}$ is the first value of a quad of values stored in element $A_{1,0}$, whereas $B_{2,4,3}$ is the fourth value of a quad of values stored in element $B_{2,4}$. In certain embodiments, processor 2100 is to convert, using convert circuits 2116A and 2116B, K (=3) quads of 8-bit floating point values from row M (=1) of the specified first source matrix (e.g., tile) 2112A and K (=3) quads of 8-bit floating point values from column N (=0) of the specified second source matrix (e.g., tile) 2112B to single-precision. In certain embodiments, processor 2100 (for example, matrix operations circuitry, e.g., as part of a matrix operations accelerator) is then to use multiply circuit 2118 to multiply the K pairs of converted first quad values from the two specified source matrices (e.g., tiles) to generate K first products, multiply the K pairs of converted second quad values from the specified source matrices (e.g., tiles) to generate K second products, multiply the K pairs of converted third quad values from the specified source matrices (e.g., tiles) to generate K third products, and multiply the K pairs of converted fourth quad values from the specified source matrices (e.g., tiles) to generate K fourth products, and then use accumulate circuit 2120 to accumulate those products separately, including a first sum of the K first products, a second sum of the K second products, a third sum of the K third products, and a fourth sum of the K fourth products, accumulate the first, second, third, and fourth sums with previous contents of the element (M,N), e.g., shown in the example here as the FP32 value from element $C(1,0)$.

20 **[0129]** In certain embodiments, processor 2100 is to use accumulate circuit 2120 to accumulate the K first products with the K second products, and separately accumulate the K third products with the K fourth products, then accumulate the two accumulated sums with previous contents of the element (M,N).

25 **[0130]** As shown, three arrows travel from each of the specified first and second source matrices (e.g., tiles), to indicate that the conversions and multiplications occur in parallel. In some embodiments, the processor responds to the decoded instruction by generating and storing results into every element of the specified destination matrix (e.g., tile) in parallel. In some embodiments, new values are generated and stored into the destination on a row-at-a-time or a column-at-a-time basis.

30 **[0131]** Disclosed embodiments improve upon alternative approaches by allowing software to perform a TDPBF8PS instruction with reduced source element sizes, which allows less memory space and less memory bandwidth to be used and improves the peak compute performance (and energy efficiency) of certain applications. Some applications, such as deep learning, can operate on input data elements with relatively few bits without losing accuracy, as long as the output elements have enough bits (e.g., more than the inputs).

35 **[0132]** **Figure 21B** is a block diagram illustrating example execution circuitry 2114 (e.g., matrix operations circuitry) to execute a TDPBF8PS instruction, according to some embodiments. Example execution circuitry 2114 includes a first data width data path (for example, 8 bits wide, e.g., according to a BF8 format) and second, wider data path width (for example, 32 bits wide, e.g., according to a full-precision format), for example, with the lines 2132 being 8 bits wide (e.g., BF8) and the lines 2134 being 32 bits wide (e.g., float32). For example, with convert BF8 to full-precision (BF8 to F32) circuits, full precision multiply (F32 MUL) circuits, and full precision addition (F32 ADD) circuits. In certain embodiments, some precision addition circuits (e.g., adder 2136) further includes bits alignment shifter, bits alignment adder, bits normalization subtractor, bits incrementor, and/or exponent logic. Systems and methods to execute a TDPBF8PS instruction are further illustrated and described, at least with respect to **Figures 22A-B, 23, and 28A-B.**

40 **[0133]** In certain embodiments, a TDPBF8PS instruction is part of a tile (e.g., AMX) architecture extension to an ISA that includes two-dimensional (2D) registers (for example, with each tile register identified as a single "tile register" (e.g., a single pointer to a single tile register), e.g., in contrast to vector (e.g., ZMM, YMM, or XMM) registers), and the ISA

may include separate instructions to load/store 2D blocks from memory (e.g., strided sets of contiguous locations), instructions to perform matrix-matrix multiplication on three registers (e.g., matrix $C_{\text{updated}} = \text{matrix } A \times \text{matrix } B + \text{matrix } C_{\text{previous}}$), and/or instructions to perform elementwise arithmetic operations on two (or three) source tiles. In one embodiment, a source matrix or source matrices are first loaded (e.g., via a host processor) into a cache (e.g., a first level (L1) data cache) and are then loaded (e.g., via execution of a tile load instruction) into a tile register (e.g., of a matrix operations accelerator) from the cache, e.g., via coherent memory interface 303 in Figure 3.

EXEMPLARY EXECUTION

[0134] Figure 22A is pseudocode illustrating exemplary execution of a TDPBF8PS instruction according to some embodiments. As shown, instruction 2201 includes an opcode 2202 (e.g. TDPBF8PS) and (e.g., tile) locations of a M by N destination matrix 2204 having single-precision elements, a M by K first source matrix 2206, and a K by N second source matrix 2208, the specified source matrices having elements comprising a quadruple of 8-bit floating-point values. Opcode 2202 (TDPBF8PS) indicates that the processor is to, as shown in pseudocode 2200, for each element (M,N) of the specified destination matrix (e.g., tile), convert K quadruples of 8-bit values from row M of the specified first source matrix (e.g., tile) and K quadruples of 8-bit values from column N of the specified second source matrix (e.g., tile) to single-precision, multiply K pairs of the converted first quad values from the two specified source matrices (e.g., tiles) to generate first products, multiply K pairs of the converted second quad values from the two specified source matrices (e.g., tiles) to generate second products, multiply K pairs of the converted third quad values from the two specified source matrices (e.g., tiles) to generate third products, multiply K pairs of the converted fourth quad values from the two specified source matrices (e.g., tiles) to generate fourth products, accumulate the first and second products to generate a first accumulated sum, separately accumulate the third and fourth products to generate a second accumulated sum, add the first and second accumulated sums to a final accumulated sum to be added with the previous contents of element (M,N). In other embodiments, the multiplications take place before the conversions.

[0135] In one embodiment, an architectural machine specific register (MSR) (e.g., as one of registers 1315 in Figure 13) (e.g., MXCSR register storing control and/or status information for an SSE register) is read (e.g., as part of the execution of an instruction), e.g., to determine exception information. DAZ may refer to a "denormals-are-zero" control (e.g., in an MSR). In certain embodiments, 8-bit precision (e.g., BF8) values are able to be processed having denormal/subnormal values.

[0136] In one embodiment, an architectural machine specific register (MSR) MXCSR register (e.g., MXCSR register storing control and/or status information for an SSE register) is not read (e.g., is not inspected and/or is not updated) (e.g., as part of the execution of an instruction). In certain embodiments, exception information for an instruction is implicit in the instruction, for example, DAZ=1 being implied for a BF8 operation (e.g., without consulting MXCSR) (e.g., for a TDPBF8PS instruction) and/or DAZ=0 being implied for a non-BF8 operation (e.g., without consulting MXCSR).

[0137] In operation, M, K, and N may be specified in one or more of several ways: as operands to the TDPBF8PS instruction (e.g., shown as tile registers "t" here), as suffixes or prefixes to the specified opcode (an asterisk is used herein as a shorthand to refer to those optional suffixes and prefixes), as part of an immediate provided with the instruction (e.g., K, M, and N each to be specified as a different (e.g., 8) bits of an (e.g., 32-bit) immediate), as part of control registers programmed by software (e.g., XTILECONFIG is a register loaded by either a matrix accelerator configuration instruction, such as TILECFG or a XRSTORE* instructions, and is stored by matrix save instruction, such as XSAVE*), or even as architectural default values.

[0138] Instruction 2201 further specifies destination matrix (e.g., tile) location 2204, first source matrix (e.g., tile) location 2206, and second source matrix (e.g., tile) location 2208. Each specified matrix (e.g., tile) locations can point to any of a memory location, a collection of vector registers, and a collection of tile registers.

[0139] Figure 22B is pseudocode 2220 for exemplary helper functions for use with a TDPBF8PS instruction, according to some embodiments. As shown, pseudocode 2220 defines a convert_bf8_to_fp32() function, a write_row_and_zero() function, a zero_upper_rows() function, and a zero_tileconfig_start() function, all of which may be used by TDPBF8PS pseudocode of Figure 22A.

[0140] Execution of a TDPBF8PS instruction is further illustrated and described with respect to Figures 21, 22A-B, 23, 28A-B, and 29A-B. Example formats of TDPBF8PS instructions is further illustrated and described with respect to Figures 24-26D.

EXEMPLARY METHOD(S) OF EXECUTION

[0141] Figure 23 is a block flow diagram illustrating a processor responding to a TDPBF8PS instruction. As shown in flow diagram 2300, at 2301, the processor is to fetch, using fetch circuitry, an instruction having fields to specify an opcode and locations of a M by N destination matrix having single-precision elements, a M by K first source matrix, and a K by N second source matrix, the specified source matrices having elements comprising a quadruple of 8-bit floating-

point values.

[0142] In embodiments that use a processor's physical register file (for example, or use one or more two-dimensional (2D) (e.g., AMX) tile registers, e.g., tile registers formed from data buffers 305 in Figure 3, which in certain embodiments are separate from any scalar and/or vector (e.g., one dimensional array) registers) to store matrices (e.g., tiles), since the destination elements are four-times as wide as the source elements, having a quadruple of 8-bit floating-point format values in the source allows efficient use, e.g., when matrices (e.g., tiles) are a collection of vector registers, of the same type of vector register, be it a 128-bit xmm register, a 256-bit ymm register, or a 512-bit zmm registers. Such efficient use can also be achieved when the matrices are stored in (e.g., AMX) tile registers. In other embodiments, a single source vector having 8-bit floating-point elements is converted into 32-bit elements stored in a destination vector having a quarter of the width of the source vector.

[0143] In certain embodiments, the specified opcode is to indicate that execution circuitry is to, for each element (M,N) of the specified destination matrix, convert K quadruples of elements from row M of the specified first source matrix and k quadruples of elements from column n of the specified second source matrix to single-precision, multiply K pairs of converted single-precision elements of a first quad position from the two specified source matrices (e.g., tiles) to generate K first products, multiple K pairs of the converted single-precision elements of a second quad position from the two specified source matrices (e.g., tiles) to generate K second products, multiply K pairs of the converted elements of a third quad position from the two specified source matrices (e.g., tiles) to generate K third products, multiply K pairs of the converted elements of a fourth quad position from the two specified source matrices (e.g., tiles) to generate K fourth products, and then accumulate the first and second products to generate a first accumulated sum and separately accumulate the third and fourth products to generate a second accumulated sum, add the first and second accumulated sums to generate a final accumulated sum to be added with previous contents of the element (m, n).

[0144] In certain embodiments, the specified opcode is to indicate execution circuitry is to cause, for each element of the first source matrix and corresponding element of the second source matrix, a conversion of the 8-bit floating-point values to single-precision values, a multiplication of converted single-precision values from first values of the quadruples together to generate a first result, a multiplication of converted single-precision values from second values of the quadruples together to generate a second result, a multiplication of converted single-precision values from third values of the quadruples together to generate a third result, a multiplication of converted single-precision values from fourth values of the quadruples together to generate a fourth result, and an accumulation of the first, second, third, and fourth results with previous contents of a corresponding element of the destination matrix.

[0145] At 2303, the processor is to decode, using decode circuitry, the fetched instruction. For example, the fetched TDPBF8PS instruction is decoded by decode circuitry such as that detailed herein. In the context of the illustrated system, decode circuitry may be that illustrated and described at least with respect to **Figures 13, 14, and 28A-B**.

[0146] At 2305, execution of the decoded instruction is scheduled (as needed), which is optional (as indicated by its dashed border) insofar as it may occur at a different time, or not at all. At 2307, the processor is to respond, using execution circuitry, to the decoded instruction as specified by the opcode.

[0147] In some embodiments, the instruction is committed or retired at 2309, which is optional (as indicated by its dashed border) insofar as it may occur at a different time, or not at all.

[0148] Example execution circuitry is further illustrated and described with respect to **Figures 3-14**. In some embodiments, execution circuitry caused execution by (e.g., offload to) a matrix operations accelerator, such as that illustrated and described as accelerator 307 (**Figure 3**). In some embodiments, execution circuitry is a matrix operations circuit, such as matrix operations circuitry 405 (**Figure 4**), 505 (**Figure 5**), or 1213 (**Figure 12**), and 1327 (**Figure 13**).

EXEMPLARY INSTRUCTION FORMAT(S)

[0149] **Figure 24** is a block diagram illustrating a format of a TDPBF8PS instruction, according to some embodiments. As shown, TDPBF8PS instruction 2400 includes fields to specify an opcode 2402 (TDPBF8PS*), which indicates that the processor is to, for each element (M,N) of the specified destination matrix, convert K quadruples of elements from row M of the specified first source matrix and K quadruples of elements from column N of the specified second source matrix to single-precision, multiply K pairs of the converted first quad values of the first quad position from the two specified source matrices (e.g., tiles) to generate K first products, multiply K pairs of the converted second quad values of the second quad position from the two specified source matrices (e.g., tiles) to generate K second products, multiply K pairs of the converted third quad values of the third quad position from the two specified source matrices (e.g., tiles) to generate K third products, multiply K pairs the converted values of the fourth quad position from the two specified source matrices (e.g., tiles) to generate K fourth products, accumulate the first and second products to generate a first accumulated sum, separately accumulate the third and fourth products to generate a second accumulated sum, add the first and second accumulated sums to a final accumulated sum to be added with the previous contents of element (m,n).

[0150] Instruction 2400 further includes destination matrix (e.g., tile) location 2404, first source matrix (e.g., tile) location 2406, and second source matrix (e.g., tile) location 2408. Each of the specified source and destination matrix locations

can be in any of a memory location, a collection of vector registers, and a collection of (e.g., AMX) tile registers.

[0151] TDPBF8PS instruction 2400 further includes several optional parameters to control the processor's behavior, including source element format 2410, k (mask control) and/or z (zeroing control) 2412, M 2414, and N 2416. In some embodiments, M and N are each any one of 4, 8, 16, and 32 (e.g., any number for either M or N, which could be 32, 64, or larger). In some embodiments, M and N are each an integer larger than or equal to 4.

[0152] Opcode 2402 is shown including an asterisk, which is to convey that additional prefixes and/or suffixes may be added to specify instruction behavior. One or more of instructions modifiers 2410, 2412, 2414, and 2416 may be specified using prefixes or suffixes to opcode 2402, e.g., a prefixes and/or suffix that indicates an instruction is to be executed with a matrix operations accelerator (e.g., including a FMA grid).

[0153] In some embodiments, one or more of optional instructions modifiers 2410, 2412, 2414, and 2416 are encoded in an immediate field (not shown) optionally included with the instruction 2400. In some embodiments, one or more of optional instructions modifiers 2410, 2412, 2414, and 2416 are specified via a configuration/status register (e.g., XTI-LECONFIG).

[0154] In some embodiments, instructions modifier 2410 includes a mask {k} 208 (e.g., writemask) and/or a zeroing control {z}, e.g., with mask {k} to control which destination elements are to be updated and/or zeroing control {z} to control whether to apply zeroing (or merging) to masked destination elements.

[0155] When any one or more of optional modifiers 2410, 2412, 2414, or 2416 are not specified by the instruction, they may use default values or implicit parameters, e.g., that are inherited from other parts of the tile architecture.

DETAILED EXEMPLARY SYSTEMS, PROCESSORS, AND EMULATION

[0156] Detailed herein are examples of hardware, software, etc. to execute the above described instructions. For example, what is described below details aspects of instruction execution including various pipeline stages such as fetch, decode, schedule, execute, retire, etc.

INSTRUCTION SETS

[0157] An instruction set may include one or more instruction formats. A given instruction format may define various fields (e.g., number of bits, location of bits) to specify, among other things, the operation to be performed (e.g., opcode) and the operand(s) on which that operation is to be performed and/or other data field(s) (e.g., mask). Some instruction formats are further broken down through the definition of instruction templates (or subformats). For example, the instruction templates of a given instruction format may be defined to have different subsets of the instruction format's fields (the included fields are typically in the same order, but at least some have different bit positions because there are less fields included) and/or defined to have a given field interpreted differently. Thus, each instruction of an ISA is expressed using a given instruction format (and, if defined, in a given one of the instruction templates of that instruction format) and includes fields for specifying the operation and the operands. For example, an exemplary ADD instruction has a specific opcode and an instruction format that includes an opcode field to specify that opcode and operand fields to select operands (source1/destination and source2); and an occurrence of this ADD instruction in an instruction stream will have specific contents in the operand fields that select specific operands. A set of SIMD extensions referred to as the Advanced Vector Extensions (AVX) (AVX1 and AVX2) and using the Vector Extensions (VEX) coding scheme has been released and/or published (e.g., see Intel® 64 and IA-32 Architectures Software Developer's Manual, September 2014; and see Intel® Advanced Vector Extensions Programming Reference, October 2014).

EXEMPLARY INSTRUCTION FORMATS

[0158] Embodiments of the instruction(s) described herein may be embodied in different formats. Additionally, exemplary systems, architectures, and pipelines are detailed below. Embodiments of the instruction(s) may be executed on such systems, architectures, and pipelines, but are not limited to those detailed.

GENERIC VECTOR FRIENDLY INSTRUCTION FORMAT

[0159] A vector friendly instruction format is an instruction format that is suited for vector instructions (e.g., there are certain fields specific to vector operations). While embodiments are described in which both vector and scalar operations are supported through the vector friendly instruction format, alternative embodiments use only vector operations the vector friendly instruction format.

[0160] **Figures 25A-25B** are block diagrams illustrating a generic vector friendly instruction format and instruction templates thereof according to embodiments. **Figure 25A** is a block diagram illustrating a generic vector friendly instruction format and class A instruction templates thereof according to embodiments; while **Figure 25B** is a block diagram

illustrating the generic vector friendly instruction format and class B instruction templates thereof according to embodiments. Specifically, a generic vector friendly instruction format 2500 for which are defined class A and class B instruction templates, both of which include no memory access 2505 instruction templates and memory access 2520 instruction templates. The term generic in the context of the vector friendly instruction format refers to the instruction format not

5 being tied to any specific instruction set.

[0161] While embodiments will be described in which the vector friendly instruction format supports the following: a 64 byte vector operand length (or size) with 32 bit (4 byte) or 64 bit (8 byte) data element widths (or sizes) (and thus, a 64 byte vector consists of either 16 doubleword-size elements or alternatively, 8 quadword-size elements); a 64 byte vector operand length (or size) with 16 bit (2 byte) or 8 bit (1 byte) data element widths (or sizes); a 32 byte vector operand length (or size) with 32 bit (4 byte), 64 bit (8 byte), 16 bit (2 byte), or 8 bit (1 byte) data element widths (or sizes); and a 16 byte vector operand length (or size) with 32 bit (4 byte), 64 bit (8 byte), 16 bit (2 byte), or 8 bit (1 byte) data element widths (or sizes); alternative embodiments may support more, less and/or different vector operand sizes (e.g., 256 byte vector operands) with more, less, or different data element widths (e.g., 128 bit (16 byte) data element widths).

10
15
20
[0162] The class A instruction templates in **Figure 25A** include: 1) within the no memory access 2505 instruction templates there is shown a no memory access, full round control type operation 2510 instruction template and a no memory access, data transform type operation 2515 instruction template; and 2) within the memory access 2520 instruction templates there is shown a memory access, temporal 2525 instruction template and a memory access, non-temporal 2530 instruction template. The class B instruction templates in **Figure 25B** include: 1) within the no memory access 2505 instruction templates there is shown a no memory access, write mask control, partial round control type operation 2512 instruction template and a no memory access, write mask control, vsize type operation 2517 instruction template; and 2) within the memory access 2520 instruction templates there is shown a memory access, write mask control 2527 instruction template.

[0163] The generic vector friendly instruction format 2500 includes the following fields listed below in the order illustrated in **Figures 25A-25B**.

25
[0164] Format field 2540 - a specific value (an instruction format identifier value) in this field uniquely identifies the vector friendly instruction format, and thus occurrences of instructions in the vector friendly instruction format in instruction streams. As such, this field is optional in the sense that it is not needed for an instruction set that has only the generic vector friendly instruction format.

[0165] Base operation field 2542 - its content distinguishes different base operations.

30
35
[0166] Register index field 2544 - its content, directly or through address generation, specifies the locations of the source and destination operands, be they in registers or in memory. These include a sufficient number of bits to select N registers from a PxQ (e.g. 32x512, 16x128, 32x1024, 64x1024) register file. While one embodiment supports up to three sources and one destination register, alternative embodiments may support more or less sources and destination registers (e.g., may support up to two sources where one of these sources also acts as the destination, may support up to three sources where one of these sources also acts as the destination, may support up to two sources and one destination).

40
[0167] Modifier field 2546 - its content distinguishes occurrences of instructions in the generic vector instruction format that specify memory access from those that do not; that is, between no memory access 2505 instruction templates and memory access 2520 instruction templates. Memory access operations read and/or write to the memory hierarchy (in some cases specifying the source and/or destination addresses using values in registers), while non-memory access operations do not (e.g., the source and destinations are registers). While in one embodiment this field also selects between three different ways to perform memory address calculations, alternative embodiments may support more, less, or different ways to perform memory address calculations.

45
[0168] Augmentation operation field 2550 - its content distinguishes which one of a variety of different operations to be performed in addition to the base operation. This field is context specific. In one embodiment, this field is divided into a class field 2568, an alpha field 2552, and a beta field 2554. The augmentation operation field 2550 allows common groups of operations to be performed in a single instruction rather than 2, 3, or 4 instructions.

[0169] Scale field 2560 - its content allows for the scaling of the index field's content for memory address generation (e.g., for address generation that uses $2^{\text{scale}} * \text{index} + \text{base}$).

50
[0170] Displacement Field 2562A- its content is used as part of memory address generation (e.g., for address generation that uses $2^{\text{scale}} * \text{index} + \text{base} + \text{displacement}$).

55
[0171] Displacement Factor Field 2562B (note that the juxtaposition of displacement field 2562A directly over displacement factor field 2562B indicates one or the other is used) - its content is used as part of address generation; it specifies a displacement factor that is to be scaled by the size of a memory access (N) - where N is the number of bytes in the memory access (e.g., for address generation that uses $2^{\text{scale}} * \text{index} + \text{base} + \text{scaled displacement}$). Redundant low-order bits are ignored and hence, the displacement factor field's content is multiplied by the memory operands total size (N) in order to generate the final displacement to be used in calculating an effective address. The value of N is determined by the processor hardware at runtime based on the full opcode field 2574 (described later herein) and the

data manipulation field 2554C. The displacement field 2562A and the displacement factor field 2562B are optional in the sense that they are not used for the no memory access 2505 instruction templates and/or different embodiments may implement only one or none of the two.

[0172] Data element width field 2564 - its content distinguishes which one of a number of data element widths is to be used (in some embodiments for all instructions; in other embodiments for only some of the instructions). This field is optional in the sense that it is not needed if only one data element width is supported and/or data element widths are supported using some aspect of the opcodes.

[0173] Write mask field 2570 - its content controls, on a per data element position basis, whether that data element position in the destination vector operand reflects the result of the base operation and augmentation operation. Class A instruction templates support merging-writemasking, while class B instruction templates support both merging- and zeroing-writemasking. When merging, vector masks allow any set of elements in the destination to be protected from updates during the execution of any operation (specified by the base operation and the augmentation operation); in other one embodiment, preserving the old value of each element of the destination where the corresponding mask bit has a 0. In contrast, when zeroing vector masks allow any set of elements in the destination to be zeroed during the execution of any operation (specified by the base operation and the augmentation operation); in one embodiment, an element of the destination is set to 0 when the corresponding mask bit has a 0 value. A subset of this functionality is the ability to control the vector length of the operation being performed (that is, the span of elements being modified, from the first to the last one); however, it is not necessary that the elements that are modified be consecutive. Thus, the write mask field 2570 allows for partial vector operations, including loads, stores, arithmetic, logical, etc. While embodiments are described in which the write mask field's 2570 content selects one of a number of write mask registers that contains the write mask to be used (and thus the write mask field's 2570 content indirectly identifies that masking to be performed), alternative embodiments instead or additional allow the mask write field's 2570 content to directly specify the masking to be performed.

[0174] Immediate field 2572 - its content allows for the specification of an immediate. This field is optional in the sense that it is not present in an implementation of the generic vector friendly format that does not support immediate and it is not present in instructions that do not use an immediate.

[0175] Class field 2568 - its content distinguishes between different classes of instructions. With reference to Figures 25A-B, the contents of this field select between class A and class B instructions. In Figures 25A-B, rounded corner squares are used to indicate a specific value is present in a field (e.g., class A 2568A and class B 2568B for the class field 2568 respectively in Figures 25A-B).

INSTRUCTION TEMPLATES OF CLASS A

[0176] In the case of the non-memory access 2505 instruction templates of class A, the alpha field 2552 is interpreted as an RS field 2552A, whose content distinguishes which one of the different augmentation operation types are to be performed (e.g., round 2552A.1 and data transform 2552A.2 are respectively specified for the no memory access, round type operation 2510 and the no memory access, data transform type operation 2515 instruction templates), while the beta field 2554 distinguishes which of the operations of the specified type is to be performed. In the no memory access 2505 instruction templates, the scale field 2560, the displacement field 2562A, and the displacement scale field 2562B are not present.

NO-MEMORY ACCESS INSTRUCTION TEMPLATES - FULL ROUND CONTROL TYPE OPERATION

[0177] In the no memory access full round control type operation 2510 instruction template, the beta field 2554 is interpreted as a round control field 2554A, whose content(s) provide static rounding. While in the described embodiments the round control field 2554A includes a suppress all floating-point exceptions (SAE) field 2556 and a round operation control field 2558, alternative embodiments may support may encode both these concepts into the same field or only have one or the other of these concepts/fields (e.g., may have only the round operation control field 2558).

[0178] SAE field 2556 - its content distinguishes whether or not to disable the exception event reporting; when the SAE field's 2556 content indicates suppression is enabled, a given instruction does not report any kind of floating-point exception flag and does not raise any floating-point exception handler.

[0179] Round operation control field 2558 - its content distinguishes which one of a group of rounding operations to perform (e.g., Round-up, Round-down, Round-towards-zero and Round-to-nearest). Thus, the round operation control field 2558 allows for the changing of the rounding mode on a per instruction basis. In one embodiment where a processor includes a control register for specifying rounding modes, the round operation control field's 2550 content overrides that register value.

No MEMORY ACCESS INSTRUCTION TEMPLATES - DATA TRANSFORM TYPE OPERATION

[0180] In the no memory access data transform type operation 2515 instruction template, the beta field 2554 is interpreted as a data transform field 2554B, whose content distinguishes which one of a number of data transforms is to be performed (e.g., no data transform, swizzle, broadcast).

[0181] In the case of a memory access 2520 instruction template of class A, the alpha field 2552 is interpreted as an eviction hint field 2552B, whose content distinguishes which one of the eviction hints is to be used (in **Figure 25A**, temporal 2552B.1 and non-temporal 2552B.2 are respectively specified for the memory access, temporal 2525 instruction template and the memory access, non-temporal 2530 instruction template), while the beta field 2554 is interpreted as a data manipulation field 2554C, whose content distinguishes which one of a number of data manipulation operations (also known as primitives) is to be performed (e.g., no manipulation; broadcast; up conversion of a source; and down conversion of a destination). The memory access 2520 instruction templates include the scale field 2560, and optionally the displacement field 2562A or the displacement scale field 2562B.

[0182] Vector memory instructions perform vector loads from and vector stores to memory, with conversion support. As with regular vector instructions, vector memory instructions transfer data from/to memory in a data element-wise fashion, with the elements that are actually transferred is dictated by the contents of the vector mask that is selected as the write mask.

MEMORY ACCESS INSTRUCTION TEMPLATES - TEMPORAL

[0183] Temporal data is data likely to be reused soon enough to benefit from caching. This is, however, a hint, and different processors may implement it in different ways, including ignoring the hint entirely.

MEMORY ACCESS INSTRUCTION TEMPLATES - NON-TEMPORAL

[0184] Non-temporal data is data unlikely to be reused soon enough to benefit from caching in the 1st-level cache and should be given priority for eviction. This is, however, a hint, and different processors may implement it in different ways, including ignoring the hint entirely.

INSTRUCTION TEMPLATES OF CLASS B

[0185] In the case of the instruction templates of class B, the alpha field 2552 is interpreted as a write mask control (Z) field 2552C, whose content distinguishes whether the write masking controlled by the write mask field 2570 should be a merging or a zeroing.

[0186] In the case of the non-memory access 2505 instruction templates of class B, part of the beta field 2554 is interpreted as an RL field 2557A, whose content distinguishes which one of the different augmentation operation types are to be performed (e.g., round 2557A.1 and vector length (VSIZE) 2557A.2 are respectively specified for the no memory access, write mask control, partial round control type operation 2512 instruction template and the no memory access, write mask control, VSIZE type operation 2517 instruction template), while the rest of the beta field 2554 distinguishes which of the operations of the specified type is to be performed. In the no memory access 2505 instruction templates, the scale field 2560, the displacement field 2562A, and the displacement scale field 2562B are not present.

[0187] In the no memory access, write mask control, partial round control type operation 2510 instruction template, the rest of the beta field 2554 is interpreted as a round operation field 2559A and exception event reporting is disabled (a given instruction does not report any kind of floating-point exception flag and does not raise any floating-point exception handler).

[0188] Round operation control field 2559A - just as round operation control field 2558, its content distinguishes which one of a group of rounding operations to perform (e.g., Round-up, Round-down, Round-towards-zero and Round-to-nearest). Thus, the round operation control field 2559A allows for the changing of the rounding mode on a per instruction basis. In one embodiment where a processor includes a control register for specifying rounding modes, the round operation control field's 2550 content overrides that register value.

[0189] In the no memory access, write mask control, VSIZE type operation 2517 instruction template, the rest of the beta field 2554 is interpreted as a vector length field 2559B, whose content distinguishes which one of a number of data vector lengths is to be performed on (e.g., 128, 256, or 512 byte).

[0190] In the case of a memory access 2520 instruction template of class B, part of the beta field 2554 is interpreted as a broadcast field 2557B, whose content distinguishes whether or not the broadcast type data manipulation operation is to be performed, while the rest of the beta field 2554 is interpreted the vector length field 2559B. The memory access 2520 instruction templates include the scale field 2560, and optionally the displacement field 2562A or the displacement scale field 2562B.

[0191] With regard to the generic vector friendly instruction format 2500, a full opcode field 2574 is shown including the format field 2540, the base operation field 2542, and the data element width field 2564. While one embodiment is shown where the full opcode field 2574 includes all of these fields, the full opcode field 2574 includes less than all of these fields in embodiments that do not support all of them. The full opcode field 2574 provides the operation code (opcode).

[0192] The augmentation operation field 2550, the data element width field 2564, and the write mask field 2570 allow these features to be specified on a per instruction basis in the generic vector friendly instruction format.

[0193] The combination of write mask field and data element width field create typed instructions in that they allow the mask to be applied based on different data element widths.

[0194] The various instruction templates found within class A and class B are beneficial in different situations. In some embodiments, different processors or different cores within a processor may support only class A, only class B, or both classes. For instance, a high performance general purpose out-of-order core intended for general-purpose computing may support only class B, a core intended primarily for graphics and/or scientific (throughput) computing may support only class A, and a core intended for both may support both (of course, a core that has some mix of templates and instructions from both classes but not all templates and instructions from both classes is within the purview of the disclosure). Also, a single processor may include multiple cores, all of which support the same class or in which different cores support different class. For instance, in a processor with separate graphics and general-purpose cores, one of the graphics cores intended primarily for graphics and/or scientific computing may support only class A, while one or more of the general-purpose cores may be high-performance general-purpose cores with out of order execution and register renaming intended for general-purpose computing that support only class B. Another processor that does not have a separate graphics core, may include one more general purpose in-order or out-of-order cores that support both class A and class B. Of course, features from one class may also be implement in the other class in different embodiments. Programs written in a high level language would be put (e.g., just in time compiled or statically compiled) into an variety of different executable forms, including: 1) a form having only instructions of the class(es) supported by the target processor for execution; or 2) a form having alternative routines written using different combinations of the instructions of all classes and having control flow code that selects the routines to execute based on the instructions supported by the processor which is currently executing the code.

EXEMPLARY SPECIFIC VECTOR FRIENDLY INSTRUCTION FORMAT

[0195] **Figure 26A** is a block diagram illustrating an exemplary specific vector friendly instruction format according to embodiments. **Figure 26A** shows a specific vector friendly instruction format 2600 that is specific in the sense that it specifies the location, size, interpretation, and order of the fields, as well as values for some of those fields. The specific vector friendly instruction format 2600 may be used to extend the x86 instruction set, and thus some of the fields are similar or the same as those used in the existing x86 instruction set and extension thereof (e.g., AVX). This format remains consistent with the prefix encoding field, real opcode byte field, MOD R/M field, SIB field, displacement field, and immediate fields of the existing x86 instruction set with extensions. The fields from **Figure 25** into which the fields from **Figure 26A** map are illustrated.

[0196] It should be understood that, although embodiments are described with reference to the specific vector friendly instruction format 2600 in the context of the generic vector friendly instruction format 2500 for illustrative purposes, the disclosure is not limited to the specific vector friendly instruction format 2600 except where claimed. For example, the generic vector friendly instruction format 2500 contemplates a variety of possible sizes for the various fields, while the specific vector friendly instruction format 2600 is shown as having fields of specific sizes. By way of specific example, while the data element width field 2564 is illustrated as a one-bit field in the specific vector friendly instruction format 2600, the disclosure is not so limited (that is, the generic vector friendly instruction format 2500 contemplates other sizes of the data element width field 2564).

[0197] The generic vector friendly instruction format 2500 includes the following fields listed below in the order illustrated in **Figure 26A**.

[0198] EVEX Prefix 2602 (Bytes 0-3)- is encoded in a four-byte form.

[0199] Format Field 2540 (EVEX Byte 0, bits [7:0]) - the first byte (EVEX Byte 0) is the format field 2540 and it contains 0x62 (the unique value used for distinguishing the vector friendly instruction format in one embodiment).

[0200] The second-fourth bytes (EVEX Bytes 1-3) include a number of bit fields providing specific capability.

[0201] REX field 2605 (EVEX Byte 1, bits [7-5]) - consists of an EVEX.R bit field (EVEX Byte 1, bit [7] - R), EVEX.X bit field (EVEX byte 1, bit [6] - X), and 2557BEX byte 1, bit [5] - B). The EVEX.R, EVEX.X, and EVEX.B bit fields provide the same functionality as the corresponding VEX bit fields, and are encoded using 1s complement form, i.e. ZMM0 is encoded as 1111B, ZMM15 is encoded as 0000B. Other fields of the instructions encode the lower three bits of the register indexes as is known in the art (rrr, xxx, and bbb), so that Rrrr, Xxxx, and Bbbb may be formed by adding EVEX.R, EVEX.X, and EVEX.B.

[0202] REX' field 2510 - this is the first part of the REX' field 2510 and is the EVEX.R' bit field (EVEX Byte 1, bit [4] - R') that is used to encode either the upper 16 or lower 16 of the extended 32 register set. In one embodiment, this bit, along with others as indicated below, is stored in bit inverted format to distinguish (in the well-known x86 32-bit mode) from the BOUND instruction, whose real opcode byte is 62, but does not accept in the MOD R/M field (described below) the value of 11 in the MOD field; alternative embodiments do not store this and the other indicated bits below in the inverted format. A value of 1 is used to encode the lower 16 registers. In other words, R'Rrrr is formed by combining EVEX.R', EVEX.R, and the other RRR from other fields.

[0203] Opcode map field 2615 (EVEX byte 1, bits [3:0] - mmmm) - its content encodes an implied leading opcode byte (0F, 0F 38, or 0F 3).

[0204] Data element width field 2564 (EVEX byte 2, bit [7] - W) - is represented by the notation EVEX.W. EVEX.W is used to define the granularity (size) of the datatype (either 32-bit data elements or 64-bit data elements).

[0205] EVEX.vvv 2620 (EVEX Byte 2, bits [6:3]-vvvv)- the role of EVEX.vvv may include the following: 1) EVEX.vvv encodes the first source register operand, specified in inverted (1s complement) form and is valid for instructions with 2 or more source operands; 2) EVEX.vvv encodes the destination register operand, specified in 1s complement form for certain vector shifts; or 3) EVEX.vvv does not encode any operand, the field is reserved and should contain 111 1b. Thus, EVEX.vvv field 2620 encodes the 4 low-order bits of the first source register specifier stored in inverted (1s complement) form. Depending on the instruction, an extra different EVEX bit field is used to extend the specifier size to 32 registers.

[0206] EVEX.U 2568 Class field (EVEX byte 2, bit [2]-U) - If EVEX.U = 0, it indicates class A or EVEX.U0; if EVEX.U = 1, it indicates class B or EVEX.U1.

[0207] Prefix encoding field 2625 (EVEX byte 2, bits [1:0]-pp) - provides additional bits for the base operation field. In addition to providing support for the legacy SSE instructions in the EVEX prefix format, this also has the benefit of compacting the STIVID prefix (rather than requiring a byte to express the SIMD prefix, the EVEX prefix requires only 2 bits). In one embodiment, to support legacy SSE instructions that use a STIVID prefix (66H, F2H, F3H) in both the legacy format and in the EVEX prefix format, these legacy SIMD prefixes are encoded into the SIMD prefix encoding field; and at runtime are expanded into the legacy SIMD prefix prior to being provided to the decoder's PLA (so the PLA can execute both the legacy and EVEX format of these legacy instructions without modification). Although newer instructions could use the EVEX prefix encoding field's content directly as an opcode extension, certain embodiments expand in a similar fashion for consistency but allow for different meanings to be specified by these legacy SIMD prefixes. An alternative embodiment may redesign the PLA to support the 2-bit SIMD prefix encodings, and thus not require the expansion.

[0208] Alpha field 2552 (EVEX byte 3, bit [7] - EH; also known as EVEX.EH, EVEX.rs, EVEX.RL, EVEX.write mask control, and EVEX.N; also illustrated with α) - as previously described, this field is context specific.

[0209] Beta field 2554 (EVEX byte 3, bits [6:4]-SSS, also known as EVEX.s₂₋₀, EVEX.r₂₋₀, EVEX.rr1, EVEX.LL0, EVEX.LLB; also illustrated with $\beta\beta\beta$) - as previously described, this field is context specific.

[0210] REX' field 2510 - this is the remainder of the REX' field and is the EVEX.V' bit field (EVEX Byte 3, bit [3] - V') that may be used to encode either the upper 16 or lower 16 of the extended 32 register set. This bit is stored in bit inverted format. A value of 1 is used to encode the lower 16 registers. In other words, V'VVVV is formed by combining EVEX.V', EVEX.vvv.

[0211] Write mask field 2570 (EVEX byte 3, bits [2:0]-kkk) - its content specifies the index of a register in the write mask registers as previously described. In one embodiment, the specific value EVEX.kkk=000 has a special behavior implying no write mask is used for the particular instruction (this may be implemented in a variety of ways including the use of a write mask hardwired to all ones or hardware that bypasses the masking hardware).

[0212] Real Opcode Field 2630 (Byte 4) is also known as the opcode byte. Part of the opcode is specified in this field.

[0213] MOD R/M Field 2640 (Byte 5) includes MOD field 2642, Reg field 2644, and R/M field 2646. As previously described, the MOD field's 2642 content distinguishes between memory access and non-memory access operations. The role of Reg field 2644 can be summarized to two situations: encoding either the destination register operand or a source register operand or be treated as an opcode extension and not used to encode any instruction operand. The role of R/M field 2646 may include the following: encoding the instruction operand that references a memory address or encoding either the destination register operand or a source register operand.

[0214] Scale, Index, Base (SIB) Byte (Byte 6) - As previously described, the content of SIB 2650 is used for memory address generation. SIB.xxx 2654 and SIB.bbb 2656 - the contents of these fields have been previously referred to with regard to the register indexes Xxxx and Bbbb.

[0215] Displacement field 2562A (Bytes 7-10) - when MOD field 2642 contains 10, bytes 7-10 are the displacement field 2562A, and it works the same as the legacy 32-bit displacement (disp32) and works at byte granularity.

[0216] Displacement factor field 2562B (Byte 7) - when MOD field 2642 contains 01, byte 7 is the displacement factor field 2562B. The location of this field is that same as that of the legacy x86 instruction set 8-bit displacement (disp8), which works at byte granularity. Since disp8 is sign extended, it can only address between -128 and +127-byte offsets;

in terms of 64-byte cache lines, disp8 uses 8 bits that can be set to only four really useful values -128, -64, 0, and 64; since a greater range is often needed, disp32 is used; however, disp32 requires 4 bytes. In contrast to disp8 and disp32, the displacement factor field 2562B is a reinterpretation of disp8; when using displacement factor field 2562B, the actual displacement is determined by the content of the displacement factor field multiplied by the size of the memory operand access (N). This type of displacement is referred to as disp8*N. This reduces the average instruction length (a single byte of used for the displacement but with a much greater range). Such compressed displacement assumes that the effective displacement is multiple of the granularity of the memory access, and hence, the redundant low-order bits of the address offset do not need to be encoded. In other words, the displacement factor field 2562B substitutes the legacy x86 instruction set 8-bit displacement. Thus, the displacement factor field 2562B is encoded the same way as an x86 instruction set 8-bit displacement (so no changes in the ModRM/SIB encoding rules) with the only exception that disp8 is overloaded to disp8*N. In other words, there are no changes in the encoding rules or encoding lengths but only in the interpretation of the displacement value by hardware (which needs to scale the displacement by the size of the memory operand to obtain a byte-wise address offset). Immediate field 2572 operates as previously described.

FULL OPCODE FIELD

[0217] Figure 26B is a block diagram illustrating the fields of the specific vector friendly instruction format 2600 that make up the full opcode field 2574 according to one embodiment. Specifically, the full opcode field 2574 includes the format field 2540, the base operation field 2542, and the data element width (W) field 2564. The base operation field 2542 includes the prefix encoding field 2625, the opcode map field 2615, and the real opcode field 2630.

REGISTER INDEX FIELD

[0218] Figure 26C is a block diagram illustrating the fields of the specific vector friendly instruction format 2600 that make up the register index field 2544 according to one embodiment. Specifically, the register index field 2544 includes the REX 2605 field, the REX' 2610 field, the MODR/M.reg field 2644, the MODR/M.r/m field 2646, the VVVV field 2620, xxx field 2654, and the bbb field 2656.

AUGMENTATION OPERATION FIELD

[0219] Figure 26D is a block diagram illustrating the fields of the specific vector friendly instruction format 2600 that make up the augmentation operation field 2550 according to one embodiment. When the class (U) field 2568 contains 0, it signifies EVEX.U0 (class A 2568A); when it contains 1, it signifies EVEX.U1 (class B 2568B). When U=0 and the MOD field 2642 contains 11 (signifying a no memory access operation), the alpha field 2552 (EVEX byte 3, bit [7] - EH) is interpreted as the rs field 2552A. When the rs field 2552A contains a 1 (round 2552A.1), the beta field 2554 (EVEX byte 3, bits [6:4]- SSS) is interpreted as the round control field 2554A. The round control field 2554A includes a one-bit SAE field 2556 and a two-bit round operation field 2558. When the rs field 2552A contains a 0 (data transform 2552A.2), the beta field 2554 (EVEX byte 3, bits [6:4]- SSS) is interpreted as a three-bit data transform field 2554B. When U=0 and the MOD field 2642 contains 00, 01, or 10 (signifying a memory access operation), the alpha field 2552 (EVEX byte 3, bit [7] - EH) is interpreted as the eviction hint (EH) field 2552B and the beta field 2554 (EVEX byte 3, bits [6:4]- SSS) is interpreted as a three-bit data manipulation field 2554C.

[0220] When U=1, the alpha field 2552 (EVEX byte 3, bit [7] - EH) is interpreted as the write mask control (Z) field 2552C. When U=1 and the MOD field 2642 contains 11 (signifying a no memory access operation), part of the beta field 2554 (EVEX byte 3, bit [4]- S₀) is interpreted as the RL field 2557A; when it contains a 1 (round 2557A. 1) the rest of the beta field 2554 (EVEX byte 3, bit [6-5]- S₂₋₁) is interpreted as the round operation field 2559A, while when the RL field 2557A contains a 0 (VSIZE 2557A.2) the rest of the beta field 2554 (EVEX byte 3, bit [6-5]- S₂₋₁) is interpreted as the vector length field 2559B (EVEX byte 3, bit [6-5]- L₁₋₀). When U=1 and the MOD field 2642 contains 00, 01, or 10 (signifying a memory access operation), the beta field 2554 (EVEX byte 3, bits [6:4]- SSS) is interpreted as the vector length field 2559B (EVEX byte 3, bit [6-5]- L₁₋₀) and the broadcast field 2557B (EVEX byte 3, bit [4]- B).

EXEMPLARY REGISTER ARCHITECTURE

[0221] Figure 27 is a block diagram of a register architecture 2700 according to one embodiment. In the embodiment illustrated, there are 32 vector registers 2710 that are 512 bits wide; these registers are referenced as zmm0 through zmm31. The lower order 256 bits of the lower 16 zmm registers are overlaid on registers ymm0-16. The lower order 128 bits of the lower 16 zmm registers (the lower order 128 bits of the ymm registers) are overlaid on registers xmm0-15. The specific vector friendly instruction format 2600 operates on these overlaid register file as illustrated in the below tables.

Adjustable Vector Length	Class	Operations	Registers
Instruction Templates that do not include the vector length field 2559B	A (Figure 25A; U=0)	2510, 2515, 2525, 2530	zmm registers (the vector length is 64 byte)
	B (Figure 25B; U=1)	2512	zmm registers (the vector length is 64 byte)
Instruction templates that do include the vector length field 2559B	B (Figure 25B; U=1)	2517, 2527	zmm, ymm, or xmm registers (the vector length is 64-byte, 32 byte, or 16 byte) depending on the vector length field 2559B

[0222] In other words, the vector length field 2559B selects between a maximum length and one or more other shorter lengths, where each such shorter length is half the length of the preceding length; and instructions templates without the vector length field 2559B operate on the maximum vector length. Further, in one embodiment, the class B instruction templates of the specific vector friendly instruction format 2600 operate on packed or scalar single/double-precision floating-point data and packed or scalar integer data. Scalar operations are operations performed on the lowest order data element position in a zmm/ymm/xmm register; the higher order data element positions are either left the same as they were prior to the instruction or zeroed depending on the embodiment.

[0223] Write mask registers 2715 - in the embodiment illustrated, there are 8 write mask registers (k0 through k7), each 64 bits in size. In an alternate embodiment, the write mask registers 2715 are 16 bits in size. As previously described, in one embodiment, the vector mask register k0 cannot be used as a write mask; when the encoding that would normally indicate k0 is used for a write mask, it selects a hardwired write mask of 0xFFFF, effectively disabling write masking for that instruction.

[0224] General-purpose registers 2725 - in the embodiment illustrated, there are sixteen 64-bit general-purpose registers that are used along with the existing x86 addressing modes to address memory operands. These registers are referenced by the names RAX, RBX, RCX, RDX, RBP, RSI, RDI, RSP, and R8 through R15.

[0225] Scalar floating-point stack register file (x87 stack) 2745, on which is aliased the MMX packed integer flat register file 2750 - in the embodiment illustrated, the x87 stack is an eight-element stack used to perform scalar floating-point operations on 32/64/80-bit floating-point data using the x87 instruction set extension; while the MMX registers are used to perform operations on 64-bit packed integer data, as well as to hold operands for some operations performed between the MMX and XMM registers.

[0226] Alternative embodiments may use wider or narrower registers. Additionally, alternative embodiments may use more, less, or different register files and registers.

EXEMPLARY CORE ARCHITECTURES, PROCESSORS, AND COMPUTER ARCHITECTURES

[0227] Processor cores may be implemented in different ways, for different purposes, and in different processors. For instance, implementations of such cores may include: 1) a general purpose in-order core intended for general-purpose computing; 2) a high-performance general purpose out-of-order core intended for general-purpose computing; 3) a special purpose core intended primarily for graphics and/or scientific (throughput) computing. Implementations of different processors may include: 1) a CPU including one or more general purpose in-order cores intended for general-purpose computing and/or one or more general purpose out-of-order cores intended for general-purpose computing; and 2) a coprocessor including one or more special purpose cores intended primarily for graphics and/or scientific (throughput). Such different processors lead to different computer system architectures, which may include: 1) the coprocessor on a separate chip from the CPU; 2) the coprocessor on a separate die in the same package as a CPU; 3) the coprocessor on the same die as a CPU (in which case, such a coprocessor is sometimes referred to as special purpose logic, such as integrated graphics and/or scientific (throughput) logic, or as special purpose cores); and 4) a system on a chip that may include on the same die the described CPU (sometimes referred to as the application core(s) or application processor(s)), the above described coprocessor, and additional functionality. Exemplary core architectures are described next, followed by descriptions of exemplary processors and computer architectures.

EXEMPLARY CORE ARCHITECTURES

IN-ORDER AND OUT-OF-ORDER CORE BLOCK DIAGRAM

[0228] Figure 28A is a block diagram illustrating both an exemplary in-order pipeline and an exemplary register

renaming, out-of-order issue/execution pipeline according to embodiments. **Figure 28B** is a block diagram illustrating both an exemplary embodiment of an in-order architecture core and an exemplary register renaming, out-of-order issue/execution architecture core to be included in a processor according to embodiments. The solid lined boxes in Figures 28A-B illustrate the in-order pipeline and in-order core, while the optional addition of the dashed lined boxes illustrates the register renaming, out-of-order issue/execution pipeline and core. Given that the in-order aspect is a subset of the out-of-order aspect, the out-of-order aspect will be described.

[0229] In **Figure 28A**, a processor pipeline 2800 includes a fetch stage 2802, a length-decode stage 2804, a decode stage 2806, an allocation stage 2808, a renaming stage 2810, a scheduling (also known as a dispatch or issue) stage 2812, a register read/memory read stage 2814, an execute stage 2816, a write back/memory write stage 2818, an exception handling stage 2822, and a commit stage 2824.

[0230] **Figure 28B** shows processor core 2890 including a front-end unit 2830 coupled to an execution engine unit 2850, and both are coupled to a memory unit 2870. The core 2890 may be a reduced instruction set computing (RISC) core, a complex instruction set computing (CISC) core, a very long instruction word (VLIW) core, or a hybrid or alternative core type. As yet another option, the core 2890 may be a special-purpose core, such as, for example, a network or communication core, compression engine, coprocessor core, general purpose computing graphics processing unit (GPGPU) core, graphics core, or the like.

[0231] The front-end unit 2830 includes a branch prediction unit 2832 coupled to an instruction cache unit 2834, which is coupled to an instruction translation lookaside buffer (TLB) 2836, which is coupled to an instruction fetch unit 2838, which is coupled to a decode unit 2840. The decode unit 2840 (or decoder) may decode instructions, and generate as an output one or more micro-operations, micro-code entry points, microinstructions, other instructions, or other control signals, which are decoded from, or which otherwise reflect, or are derived from, the original instructions. The decode unit 2840 may be implemented using various different mechanisms. Examples of suitable mechanisms include, but are not limited to, look-up tables, hardware implementations, programmable logic arrays (PLAs), microcode read only memories (ROMs), etc. In one embodiment, the core 2890 includes a microcode ROM or other medium that stores microcode for certain macroinstructions (e.g., in decode unit 2840 or otherwise within the front-end unit 2830). The decode unit 2840 is coupled to a rename/allocator unit 2852 in the execution engine unit 2850.

[0232] The execution engine unit 2850 includes the rename/allocator unit 2852 coupled to a retirement unit 2854 and a set of one or more scheduler unit(s) 2856. The scheduler unit(s) 2856 represents any number of different schedulers, including reservations stations, central instruction window, etc. The scheduler unit(s) 2856 is coupled to the physical register file(s) unit(s) 2858. Each of the physical register file(s) units 2858 represents one or more physical register files, different ones of which store one or more different data types, such as scalar integer, scalar floating-point, packed integer, packed floating-point, vector integer, vector floating-point, status (e.g., an instruction pointer that is the address of the next instruction to be executed), etc. In one embodiment, the physical register file(s) unit 2858 comprises a vector registers unit, a write mask registers unit, and a scalar registers unit. These register units may provide architectural vector registers, vector mask registers, and general-purpose registers. The physical register file(s) unit(s) 2858 is overlapped by the retirement unit 2854 to illustrate various ways in which register renaming and out-of-order execution may be implemented (e.g., using a reorder buffer(s) and a retirement register file(s); using a future file(s), a history buffer(s), and a retirement register file(s); using a register maps and a pool of registers; etc.). The retirement unit 2854 and the physical register file(s) unit(s) 2858 are coupled to the execution cluster(s) 2860. The execution cluster(s) 2860 includes a set of one or more execution units 2862 and a set of one or more memory access units 2864. The execution units 2862 may perform various operations (e.g., shifts, addition, subtraction, multiplication) and on various types of data (e.g., scalar floating-point, packed integer, packed floating-point, vector integer, vector floating-point). While some embodiments may include a number of execution units dedicated to specific functions or sets of functions, other embodiments may include only one execution unit or multiple execution units that all perform all functions. The scheduler unit(s) 2856, physical register file(s) unit(s) 2858, and execution cluster(s) 2860 are shown as being possibly plural because certain embodiments create separate pipelines for certain types of data/operations (e.g., a scalar integer pipeline, a scalar floating-point/packed integer/packed floating-point/vector integer/vector floating-point pipeline, and/or a memory access pipeline that each have their own scheduler unit, physical register file(s) unit, and/or execution cluster - and in the case of a separate memory access pipeline, certain embodiments are implemented in which only the execution cluster of this pipeline has the memory access unit(s) 2864). It should also be understood that where separate pipelines are used, one or more of these pipelines may be out-of-order issue/execution and the rest in-order.

[0233] The set of memory access units 2864 is coupled to the memory unit 2870, which includes a data TLB unit 2872 coupled to a data cache unit 2874 coupled to a level 2 (L2) cache unit 2876. In one exemplary embodiment, the memory access units 2864 may include a load unit, a store address unit, and a store data unit, each of which is coupled to the data TLB unit 2872 in the memory unit 2870. The instruction cache unit 2834 is further coupled to a level 2 (L2) cache unit 2876 in the memory unit 2870. The L2 cache unit 2876 is coupled to one or more other levels of cache and eventually to a main memory.

[0234] By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may im-

plement the pipeline 2800 as follows: 1) the instruction fetch 2838 performs the fetch and length decoding stages 2802 and 2804; 2) the decode unit 2840 performs the decode stage 2806; 3) the rename/allocator unit 2852 performs the allocation stage 2808 and renaming stage 2810; 4) the scheduler unit(s) 2856 performs the schedule stage 2812; 5) the physical register file(s) unit(s) 2858 and the memory unit 2870 perform the register read/memory read stage 2814; the execution cluster 2860 perform the execute stage 2816; 6) the memory unit 2870 and the physical register file(s) unit(s) 2858 perform the write back/memory write stage 2818; 7) various units may be involved in the exception handling stage 2822; and 8) the retirement unit 2854 and the physical register file(s) unit(s) 2858 perform the commit stage 2824.

[0235] The core 2890 may support one or more instructions sets (e.g., the x86 instruction set (with some extensions that have been added with newer versions); the MIPS instruction set of MIPS Technologies of Sunnyvale, CA; the ARM instruction set (with optional additional extensions such as NEON) of ARM Holdings of Sunnyvale, CA), including the instruction(s) described herein. In one embodiment, the core 2890 includes logic to support a packed data instruction set extension (e.g., AVX1, AVX2), thereby allowing the operations used by many multimedia applications to be performed using packed data.

[0236] It should be understood that the core may support multithreading (executing two or more parallel sets of operations or threads), and may do so in a variety of ways including time sliced multithreading, simultaneous multithreading (where a single physical core provides a logical core for each of the threads that physical core is simultaneously multithreading), or a combination thereof (e.g., time sliced fetching and decoding and simultaneous multithreading thereafter such as in the Intel® Hyperthreading technology).

[0237] While register renaming is described in the context of out-of-order execution, it should be understood that register renaming may be used in an in-order architecture. While the illustrated embodiment of the processor also includes separate instruction and data cache units 2834/2874 and a shared L2 cache unit 2876, alternative embodiments may have a single internal cache for both instructions and data, such as, for example, a Level 1 (L1) internal cache, or multiple levels of internal cache. In some embodiments, the system may include a combination of an internal cache and an external cache that is external to the core and/or the processor. Alternatively, all of the cache may be external to the core and/or the processor.

SPECIFIC EXEMPLARY IN-ORDER CORE ARCHITECTURE

[0238] Figures 29A-B illustrate a block diagram of a more specific exemplary in-order core architecture, which core would be one of several logic blocks (including other cores of the same type and/or different types) in a chip. The logic blocks communicate through a high-bandwidth interconnect network (e.g., a ring network) with some fixed function logic, memory I/O interfaces, and other necessary I/O logic, depending on the application.

[0239] **Figure 29A** is a block diagram of a single processor core, along with its connection to the on-die interconnect network 2902 and with its local subset of the Level 2 (L2) cache 2904, according to embodiments. In one embodiment, an instruction decoder 2900 supports the x86 instruction set with a packed data instruction set extension. An L1 cache 2906 allows low-latency accesses to cache memory into the scalar and vector units. While in one embodiment (to simplify the design), a scalar unit 2908 and a vector unit 2910 use separate register sets (respectively, scalar registers 2912 and vector registers 2914) and data transferred between them is written to memory and then read back in from a level 1 (L1) cache 2906, alternative embodiments may use a different approach (e.g., use a single register set or include a communication path that allow data to be transferred between the two register files without being written and read back).

[0240] The local subset of the L2 cache 2904 is part of a global L2 cache that is divided into separate local subsets, one per processor core. Each processor core has a direct access path to its own local subset of the L2 cache 2904. Data read by a processor core is stored in its L2 cache subset 2904 and can be accessed quickly, in parallel with other processor cores accessing their own local L2 cache subsets. Data written by a processor core is stored in its own L2 cache subset 2904 and is flushed from other subsets, if necessary. The ring network ensures coherency for shared data. The ring network is bi-directional to allow agents such as processor cores, L2 caches and other logic blocks to communicate with each other within the chip. Each ring data-path is 1012-bits wide per direction.

[0241] **Figure 29B** is an expanded view of part of the processor core in **Figure 29A** according to embodiments. **Figure 29B** includes an L1 data cache 2906A part of the L1 cache 2904, as well as more detail regarding the vector unit 2910 and the vector registers 2914. Specifically, the vector unit 2910 is a 16-wide vector processing unit (VPU) (see the 16-wide ALU 2928), which executes one or more of integer, single-precision float, and double-precision float instructions. The VPU supports swizzling the register inputs with swizzle unit 2920, numeric conversion with numeric convert units 2922A-B, and replication with replication unit 2924 on the memory input. Write mask registers 2926 allow predicating resulting vector writes.

[0242] **Figure 30** is a block diagram of a processor 3000 that may have more than one core, may have an integrated memory controller, and may have integrated graphics according to embodiments. The solid lined boxes in **Figure 30** illustrate a processor 3000 with a single core 3002A, a system agent 3010, a set of one or more bus controller units 3016, while the optional addition of the dashed lined boxes illustrates an alternative processor 3000 with multiple cores

3002A-N, a set of one or more integrated memory controller unit(s) 3014 in the system agent unit 3010, and special purpose logic 3008.

[0243] Thus, different implementations of the processor 3000 may include: 1) a CPU with the special purpose logic 3008 being integrated graphics and/or scientific (throughput) logic (which may include one or more cores), and the cores 3002A-N being one or more general purpose cores (e.g., general purpose in-order cores, general purpose out-of-order cores, a combination of the two); 2) a coprocessor with the cores 3002A-N being a large number of special purpose cores intended primarily for graphics and/or scientific (throughput); and 3) a coprocessor with the cores 3002A-N being a large number of general purpose in-order cores. Thus, the processor 3000 may be a general-purpose processor, coprocessor or special-purpose processor, such as, for example, a network or communication processor, compression engine, graphics processor, GPGPU (general purpose graphics processing unit), a high-throughput many integrated core (MIC) coprocessor (including 30 or more cores), embedded processor, or the like. The processor may be implemented on one or more chips. The processor 3000 may be a part of and/or may be implemented on one or more substrates using any of a number of process technologies, such as, for example, BiCMOS, CMOS, or NMOS.

[0244] The memory hierarchy includes one or more levels of cache within the cores, a set of one or more shared cache units 3006, and external memory (not shown) coupled to the set of integrated memory controller units 3014. The set of shared cache units 3006 may include one or more mid-level caches, such as level 2 (L2), level 3 (L3), level 4 (L4), or other levels of cache, a last level cache (LLC), and/or combinations thereof. While in one embodiment a ring-based interconnect unit 3012 interconnects the special purpose logic 3008 (integrated graphics logic is an example of and is also referred to herein as special purpose logic), the set of shared cache units 3006, and the system agent unit 3010/integrated memory controller unit(s) 3014, alternative embodiments may use any number of well-known techniques for interconnecting such units. In one embodiment, coherency is maintained between one or more cache units 3006 and cores 3002A-N.

[0245] In some embodiments, one or more of the cores 3002A-N are capable of multi-threading. The system agent 3010 includes those components coordinating and operating cores 3002A-N. The system agent unit 3010 may include for example a power control unit (PCU) and a display unit. The PCU may be or include logic and components needed for regulating the power state of the cores 3002A-N and the special purpose logic 3008. The display unit is for driving one or more externally connected displays.

[0246] The cores 3002A-N may be homogenous or heterogeneous in terms of architecture instruction set; that is, two or more of the cores 3002A-N may be capable of execution the same instruction set, while others may be capable of executing only a subset of that instruction set or a different instruction set.

EXEMPLARY COMPUTER ARCHITECTURES

[0247] **Figures 31-34** are block diagrams of exemplary computer architectures. Other system designs and configurations known in the arts for laptops, desktops, handheld PCs, personal digital assistants, engineering workstations, servers, network devices, network hubs, switches, embedded processors, digital signal processors (DSPs), graphics devices, video game devices, set-top boxes, micro controllers, cell phones, portable media players, hand held devices, and various other electronic devices, are also suitable. In general, a huge variety of systems or electronic devices capable of incorporating a processor and/or other execution logic as disclosed herein are generally suitable.

[0248] Referring now to **Figure 31**, shown is a block diagram of a system 3100 in accordance with one embodiment of the present disclosure. The system 3100 may include one or more processors 3110, 3115, which are coupled to a controller hub 3120. In one embodiment the controller hub 3120 includes a graphics memory controller hub (GMCH) 3190 and an Input/Output Hub (IOH) 3150 (which may be on separate chips); the GMCH 3190 includes memory and graphics controllers to which are coupled memory 3140 and a coprocessor 3145; the IOH 3150 couples input/output (I/O) devices 3160 to the GMCH 3190. Alternatively, one or both of the memory and graphics controllers are integrated within the processor (as described herein), the memory 3140 and the coprocessor 3145 are coupled directly to the processor 3110, and the controller hub 3120 in a single chip with the IOH 3150. Memory 3140 may include matrix acceleration code 3140A, for example, that stores code that when executed causes a processor to perform any method of this disclosure.

[0249] The optional nature of additional processors 3115 is denoted in **Figure 31** with broken lines. Each processor 3110, 3115 may include one or more of the processing cores described herein and may be some version of the processor 3000.

[0250] The memory 3140 may be, for example, dynamic random-access memory (DRAM), phase change memory (PCM), or a combination of the two. For at least one embodiment, the controller hub 3120 communicates with the processor(s) 3110, 3115 via a multi-drop bus, such as a frontside bus (FSB), point-to-point interface such as QuickPath Interconnect (QPI), or similar connection 3195.

[0251] In one embodiment, the coprocessor 3145 is a special-purpose processor, such as, for example, a high-throughput MIC processor, a network or communication processor, compression engine, graphics processor, GPGPU,

embedded processor, or the like. In one embodiment, controller hub 3120 may include an integrated graphics accelerator.

[0252] There can be a variety of differences between the physical resources 3110, 3115 in terms of a spectrum of metrics of merit including architectural, microarchitectural, thermal, power consumption characteristics, and the like.

[0253] In one embodiment, the processor 3110 executes instructions that control data processing operations of a general type. Embedded within the instructions may be coprocessor instructions. The processor 3110 recognizes these coprocessor instructions as being of a type that should be executed by the attached coprocessor 3145. Accordingly, the processor 3110 issues these coprocessor instructions (or control signals representing coprocessor instructions) on a coprocessor bus or other interconnect, to coprocessor 3145. Coprocessor(s) 3145 accept and execute the received coprocessor instructions.

[0254] Referring now to **Figure 32**, shown is a block diagram of a first more specific exemplary system 3200 in accordance with an embodiment of the present disclosure. As shown in **Figure 32**, multiprocessor system 3200 is a point-to-point interconnect system, and includes a first processor 3270 and a second processor 3280 coupled via a point-to-point interconnect 3250. Each of processors 3270 and 3280 may be some version of the processor 3000. In one embodiment, processors 3270 and 3280 are respectively processors 3110 and 3115, while coprocessor 3238 is coprocessor 3145. In another embodiment, processors 3270 and 3280 are respectively processor 3110 coprocessor 3145.

[0255] Processors 3270 and 3280 are shown including integrated memory controller (IMC) units 3272 and 3282, respectively. Processor 3270 also includes as part of its bus controller units point-to-point (P-P) interfaces 3276 and 3278; similarly, second processor 3280 includes P-P interfaces 3286 and 3288. Processors 3270, 3280 may exchange information via a point-to-point (P-P) interface 3250 using P-P interface circuits 3278, 3288. As shown in **Figure 32**, IMCs 3272 and 3282 couple the processors to respective memories, namely a memory 3232 and a memory 3234, which may be portions of main memory locally attached to the respective processors.

[0256] Processors 3270, 3280 may each exchange information with a chipset 3290 via individual P-P interfaces 3252, 3254 using point to point interface circuits 3276, 3294, 3286, 3298. Chipset 3290 may optionally exchange information with the coprocessor 3238 via a high-performance interface 3292. In one embodiment, the coprocessor 3238 is a special-purpose processor, such as, for example, a high-throughput MIC processor, a network or communication processor, compression engine, graphics processor, GPGPU, embedded processor, or the like.

[0257] A shared cache (not shown) may be included in either processor or outside of both processors yet connected with the processors via P-P interconnect, such that either or both processors' local cache information may be stored in the shared cache if a processor is placed into a low power mode.

[0258] Chipset 3290 may be coupled to a first bus 3216 via an interface 3296. In one embodiment, first bus 3216 may be a Peripheral Component Interconnect (PCI) bus, or a bus such as a PCI Express bus or another third generation I/O interconnect bus, although the scope of the present disclosure is not so limited.

[0259] As shown in **Figure 32**, various I/O devices 3214 may be coupled to first bus 3216, along with a bus bridge 3218 which couples first bus 3216 to a second bus 3220. In one embodiment, one or more additional processor(s) 3215, such as coprocessors, high-throughput MIC processors, GPGPU's, accelerators (such as, e.g., graphics accelerators or digital signal processing (DSP) units), field programmable gate arrays, or any other processor, are coupled to first bus 3216. In one embodiment, second bus 3220 may be a low pin count (LPC) bus. Various devices may be coupled to a second bus 3220 including, for example, a keyboard and/or mouse 3222, communication devices 3227 and a storage unit 3228 such as a disk drive or other mass storage device which may include instructions/code and data 3230, in one embodiment. Further, an audio I/O 3224 may be coupled to the second bus 3220. Note that other architectures are possible. For example, instead of the point-to-point architecture of **Figure 32**, a system may implement a multi-drop bus or other such architecture.

[0260] Referring now to **Figure 33**, shown is a block diagram of a second more specific exemplary system 3300 in accordance with an embodiment of the present disclosure. Like elements in **Figures 32 and 33** bear like reference numerals, and certain aspects of **Figure 32** have been omitted from **Figure 33** in order to avoid obscuring other aspects of **Figure 33**.

[0261] **Figure 33** illustrates that the processors 3270, 3280 may include integrated memory and I/O control logic ("CL") 3372 and 3382, respectively. Thus, the CL 3372, 3382 include integrated memory controller units and include I/O control logic. **Figure 33** illustrates that not only are the memories 3232, 3234 coupled to the CL 3372, 3382, but also that I/O devices 3314 are also coupled to the control logic 3372, 3382. Legacy I/O devices 3315 are coupled to the chipset 3290.

[0262] Referring now to **Figure 34**, shown is a block diagram of a SoC 3400 in accordance with an embodiment of the present disclosure. Similar elements in **Figure 30** bear like reference numerals. Also, dashed lined boxes are optional features on more advanced SoCs. In **Figure 34**, an interconnect unit(s) 3402 is coupled to: an application processor 3410 which includes a set of one or more cores 3002A-N, which include cache units 3004A-N, and shared cache unit(s) 3006; a system agent unit 3010; a bus controller unit(s) 3016; an integrated memory controller unit(s) 3014; a set or one or more coprocessors 3420 which may include integrated graphics logic, an image processor, an audio processor, and a video processor; an static random access memory (SRAM) unit 3430; a direct memory access (DMA) unit 3432; and

a display unit 3440 for coupling to one or more external displays. In one embodiment, the coprocessor(s) 3420 include a special-purpose processor, such as, for example, a network or communication processor, compression engine, GPG-PU, a high-throughput MIC processor, embedded processor, or the like.

[0263] Embodiments of the mechanisms disclosed herein may be implemented in hardware, software, firmware, or a combination of such implementation approaches. Embodiments may be implemented as computer programs or program code executing on programmable systems comprising at least one processor, a storage system (including volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device.

[0264] Program code, such as code 3230 illustrated in **Figure 32**, may be applied to input instructions to perform the functions described herein and generate output information. The output information may be applied to one or more output devices, in known fashion. For purposes of this application, a processing system includes any system that has a processor, such as, for example; a digital signal processor (DSP), a microcontroller, an application specific integrated circuit (ASIC), or a microprocessor.

[0265] The program code may be implemented in a high level procedural or object-oriented programming language to communicate with a processing system. The program code may also be implemented in assembly or machine language, if desired. In fact, the mechanisms described herein are not limited in scope to any particular programming language. In any case, the language may be a compiled or interpreted language.

[0266] One or more aspects of at least one embodiment may be implemented by representative instructions stored on a machine-readable medium which represents various logic within the processor, which when read by a machine causes the machine to fabricate logic to perform the techniques described herein. Such representations, known as "IP cores" may be stored on a tangible, machine readable medium and supplied to various customers or manufacturing facilities to load into the fabrication machines that actually make the logic or processor.

[0267] Such machine-readable storage media may include, without limitation, non-transitory, tangible arrangements of articles manufactured or formed by a machine or device, including storage media such as hard disks, any other type of disk including floppy disks, optical disks, compact disk read-only memories (CD-ROMs), compact disk rewritable's (CD-RWs), and magneto-optical disks, semiconductor devices such as read-only memories (ROMs), random access memories (RAMs) such as dynamic random access memories (DRAMs), static random access memories (SRAMs), erasable programmable read-only memories (EPROMs), flash memories, electrically erasable programmable read-only memories (EEPROMs), phase change memory (PCM), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

[0268] Accordingly, embodiments also include non-transitory, tangible machine-readable media containing instructions or containing design data, such as Hardware Description Language (HDL), which defines structures, circuits, apparatuses, processors and/or system features described herein. Such embodiments may also be referred to as program products.

EMULATION (INCLUDING BINARY TRANSLATION, CODE MORPHING, ETC.)

[0269] In some cases, an instruction converter may be used to convert an instruction from a source instruction set to a target instruction set. For example, the instruction converter may translate (e.g., using static binary translation, dynamic binary translation including dynamic compilation), morph, emulate, or otherwise convert an instruction to one or more other instructions to be processed by the core. The instruction converter may be implemented in software, hardware, firmware, or a combination thereof. The instruction converter may be on processor, off processor, or part on and part off processor.

[0270] **Figure 35** is a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction set according to embodiments. In the illustrated embodiment, the instruction converter is a software instruction converter, although alternatively the instruction converter may be implemented in software, firmware, hardware, or various combinations thereof. **Figure 35** shows a program in a high-level language 3502 may be compiled using an x86 compiler 3504 to generate x86 binary code 3506 that may be natively executed by a processor with at least one x86 instruction set core 3516. The processor with at least one x86 instruction set core 3516 represents any processor that can perform substantially the same functions as an Intel® processor with at least one x86 instruction set core by compatibly executing or otherwise processing (1) a substantial portion of the instruction set of the Intel® x86 instruction set core or (2) object code versions of applications or other software targeted to run on an Intel® processor with at least one x86 instruction set core, in order to achieve substantially the same result as an Intel® processor with at least one x86 instruction set core. The x86 compiler 3504 represents a compiler that is operable to generate x86 binary code 3506 (e.g., object code) that can, with or without additional linkage processing, be executed on the processor with at least one x86 instruction set core 3516. Similarly, **Figure 35** shows the program in the high level language 3502 may be compiled using an alternative instruction set compiler 3508 to generate alternative instruction set binary code 3510 that may be natively executed by a processor without at least one x86 instruction set core 3514 (e.g., a processor with cores that execute the MIPS instruction set of MIPS Technologies

of Sunnyvale, CA and/or that execute the ARM instruction set of ARM Holdings of Sunnyvale, CA). The instruction converter 3512 is used to convert the x86 binary code 3506 into code that may be natively executed by the processor without an x86 instruction set core 3514. This converted code is not likely to be the same as the alternative instruction set binary code 3510 because an instruction converter capable of this is difficult to make; however, the converted code will accomplish the general operation and be made up of instructions from the alternative instruction set. Thus, the instruction converter 3512 represents software, firmware, hardware, or a combination thereof that, through emulation, simulation or any other process, allows a processor or other electronic device that does not have an x86 instruction set processor or core to execute the x86 binary code 3506.

FURTHER EXAMPLES

[0271] At least some embodiments of the disclosed technologies can be described in view of the following examples:

Example 1. An apparatus comprising:

fetch circuitry to fetch a single instruction having fields to specify an opcode and locations of a M by N destination matrix having single-precision elements, an M by K first source matrix, and a K by N second source matrix, the source matrices having elements that each comprise a quadruple of 8-bit floating-point values, the opcode to indicate execution circuitry is to cause, for each element of the first source matrix and corresponding element of the second source matrix, a conversion of the 8-bit floating-point values to single-precision values, a multiplication of converted single-precision values from first values of the quadruples together to generate a first result, a multiplication of converted single-precision values from second values of the quadruples together to generate a second result, a multiplication of converted single-precision values from third values of the quadruples together to generate a third result, a multiplication of converted single-precision values from fourth values of the quadruples together to generate a fourth result, and an accumulation of the first, second, third, and fourth results with previous contents of a corresponding element of the destination matrix;

decode circuitry to decode the fetched instruction; and

the execution circuitry to respond to the decoded instruction as specified by the opcode.

Example 2. The apparatus of example 1, wherein the 8-bit floating-point format is specified by the opcode of the single instruction.

Example 3. The apparatus of example 1, wherein M, N, and K are specified by the single instruction.

Example 4. The apparatus of example 1, where the execution circuitry is to cause a matrix operations accelerator to perform at least the multiplications and the accumulation.

Example 5. The apparatus of example 4, wherein M, N, and K are specified by a configuration of the matrix operations accelerator to be programmed by execution of a matrix accelerator configuration instruction before executing the single instruction.

Example 6. The apparatus of example 1, wherein the execution circuitry is further to cause saturation of execution results, as necessary.

Example 7. The apparatus of example 1, wherein the single instruction is further to specify a writemask comprising M x N bits, each bit to control whether to mask a corresponding element of the destination matrix.

Example 8. The apparatus of example 1, wherein the execution circuitry is further to generate a fault when a fault condition occurs, the fault condition selectable from:

the destination matrix having a fewer number of rows than a number of rows of the first source matrix; and

the destination matrix having a fewer number of columns than a number of columns of the second source matrix.

Example 9. A method comprising:

fetching, by fetch circuitry of a processor, a single instruction having fields to specify an opcode and locations of a M by N destination matrix having single-precision elements, an M by K first source matrix, and a K by N second source matrix, the source matrices having elements that each comprise a quadruple of 8-bit floating-

point values, the opcode to indicate execution circuitry is to cause, for each element of the first source matrix and corresponding element of the second source matrix, a conversion of the 8-bit floating-point values to single-precision values, a multiplication of converted single-precision values from first values of the quadruples together to generate a first result, a multiplication of converted single-precision values from second values of the quadruples together to generate a second result, a multiplication of converted single-precision values from third values of the quadruples together to generate a third result, a multiplication of converted single-precision values from fourth values of the quadruples together to generate a fourth result, and an accumulation of the first, second, third, and fourth results with previous contents of a corresponding element of the destination matrix; decoding, by decode circuitry of the processor, the fetched instruction into a decoded single instruction; and executing, by the execution circuitry of the processor, the decoded single instruction according to the opcode.

Example 10. The method of example 9, wherein the 8-bit floating-point format is specified by the opcode of the single instruction.

Example 11. The method of example 9, wherein M, N, and K are specified by the single instruction.

Example 12. The method of example 9, where the execution circuitry causes a matrix operations accelerator to perform at least the multiplications and the accumulation.

Example 13. The method of example 12, further comprising executing, by the execution circuitry of the processor before executing the single instruction, a matrix accelerator configuration instruction that programs a configuration of the matrix operations accelerator specifying M, N, and K.

Example 14. The method of example 9, wherein the executing comprises saturating the execution results.

Example 15. The method of example 9, wherein the single instruction further specifies a writemask comprising M x N bits, each bit controlling whether to mask a corresponding element of the destination matrix.

Example 16. The method of example 9, wherein the executing generates a fault when a fault condition occurs, the fault condition selectable from:

the destination matrix having a fewer number of rows than a number of rows of the first source matrix; and the destination matrix having a fewer number of columns than a number of columns of the second source matrix.

Example 17. A non-transitory machine readable medium that stores program code that when executed by a machine causes the machine to perform a method comprising:

fetching, by fetch circuitry of a processor, a single instruction having fields to specify an opcode and locations of a M by N destination matrix having single-precision elements, an M by K first source matrix, and a K by N second source matrix, the source matrices having elements that each comprise a quadruple of 8-bit floating-point values, the opcode to indicate execution circuitry is to cause, for each element of the first source matrix and corresponding element of the second source matrix, a conversion of the 8-bit floating-point values to single-precision values, a multiplication of converted single-precision values from first values of the quadruples together to generate a first result, a multiplication of converted single-precision values from second values of the quadruples together to generate a second result, a multiplication of converted single-precision values from third values of the quadruples together to generate a third result, a multiplication of converted single-precision values from fourth values of the quadruples together to generate a fourth result, and an accumulation of the first, second, third, and fourth results with previous contents of a corresponding element of the destination matrix; decoding, by decode circuitry of the processor, the fetched instruction into a decoded single instruction; and executing, by the execution circuitry of the processor, the decoded single instruction according to the opcode.

Example 18. The non-transitory machine readable medium of example 17, wherein the 8-bit floating-point format is specified by the opcode of the single instruction.

Example 19. The non-transitory machine readable medium of example 17, wherein M, N, and K are specified by the single instruction.

Example 20. The non-transitory machine readable medium of example 17, where the executing comprises the

execution circuitry causing a matrix operations accelerator to perform at least the multiplications and the accumulation.

5 Example 21. The non-transitory machine readable medium of example 20, wherein the method further comprises executing, by the execution circuitry of the processor before executing the single instruction, a matrix accelerator configuration instruction that programs a configuration of the matrix operations accelerator specifying M, N, and K.

10 Example 22. The non-transitory machine readable medium of example 17, wherein the executing comprises saturating the execution results.

Example 23. The non-transitory machine readable medium of example 17, wherein the single instruction further specifies a writemask comprising M x N bits, each bit controlling whether to mask a corresponding element of the destination matrix.

15 Example 24. The non-transitory machine readable medium of example 17, wherein the executing generates a fault when a fault condition occurs, the fault condition selectable from:

the destination matrix having a fewer number of rows than a number of rows of the first source matrix; and
the destination matrix having a fewer number of columns than a number of columns of the second source matrix.

Claims

25 1. A processing unit comprising:
circuitry to perform operations corresponding to an instruction, the instruction to specify a first matrix having M rows by 4*K columns of 8-bit floating-point data elements, a second matrix having 4*K rows by N columns of 8-bit floating-point data elements, and a third matrix having M rows by N columns of 32-bit single precision floating-point data elements, the operations including to:
for each row m of the M rows of the first matrix, and for each column n of the N columns of the second matrix:

30 convert 4*K 8-bit floating-point data elements of the row m of the first matrix to 4*K corresponding higher precision floating-point data elements having a higher precision than an 8-bit floating-point data element, and convert 4*K 8-bit floating-point data elements of the column n of the second matrix to 4*K corresponding higher precision floating-point data elements having a higher precision than the 8-bit floating-point data element;
35 multiply the 4*K higher precision floating-point data elements corresponding to the row m of the first matrix with corresponding ones of the 4*K higher precision floating-point data elements corresponding to the column n of the second matrix to generate 4*K products;
accumulate the 4*K products with a 32-bit single precision floating-point data element corresponding to a row m of the M rows, and a column n of the N columns, of the third matrix, to generate a result 32-bit single precision
40 floating-point data element; and
store the result 32-bit single precision floating-point data element at the row m and the column n of the third matrix.

2. The processing unit of claim 1, wherein the M rows of the first matrix is 64 rows.

45 3. The processing unit of claim 2, wherein the N columns of the second matrix is more than 64 columns.

4. The processing unit of any of claims 1-3, wherein the instruction has a value to specify the N columns of the second matrix.

50 5. The processing unit of any of claims 1-4, wherein the first matrix is either to be stored in a plurality of registers of the processing unit or in memory, and wherein the second matrix is to be stored in memory.

6. The processing unit of any of claims 1-5, wherein the instruction has a source element format field.

55 7. The processing unit of any of claims 1-6, comprising:

a plurality of memory controllers;
a level-two, L2, cache memory coupled to the plurality of memory controllers; and

a plurality of cores coupled to the plurality of memory controllers, and coupled to the L2 cache memory, the plurality of cores comprising the circuitry.

- 5
8. The processing unit of claim 7, further comprising one or more of an interconnect interface coupled to the plurality of cores, and a bus controller interface coupled to the plurality of cores.
9. The processing unit of claim 7 or 8, wherein the plurality of cores comprises graphics cores.
- 10
10. The processing unit of any of claims 7-9, wherein the plurality of cores comprises heterogeneous cores.
11. The processing unit of any of claims 7-10, wherein the plurality of cores comprises heterogeneous graphics cores.
12. The processing unit of any of claims 7-11, further comprising an instruction converter to convert the instruction into one or more instructions of a different instruction set executable by the plurality of cores.
- 15
13. The processing unit of any of claims 1-12, wherein the 8-bit floating-point data elements each have five exponent bits and two explicit mantissa bits.
14. The processing unit of any of claims 1-13, wherein the higher precision floating-point data elements comprise 32-bits.
- 20
15. A method comprising:
 processing data with a processing unit, including performing operations corresponding to an instruction, the instruction to specify a first matrix having M rows by 4*K columns of 8-bit floating-point data elements, a second matrix having 4*K rows by N columns of 8-bit floating-point data elements, and a third matrix having M rows by N columns of 32-bit single precision floating-point data elements, the operations including to:
 25 for each row m of the M rows of the first matrix, and for each column n of the N columns of the second matrix:
- 30 convert 4*K 8-bit floating-point data elements of the row m of the first matrix to 4*K corresponding higher precision floating-point data elements having a higher precision than an 8-bit floating-point data element, and convert 4*K 8-bit floating-point data elements of the column n of the second matrix to 4*K corresponding higher precision floating-point data elements having a higher precision than the 8-bit floating-point data element;
 multiply the 4*K higher precision floating-point data elements corresponding to the row m of the first matrix with corresponding ones of the 4*K higher precision floating-point data elements corresponding to the column n of the second matrix to generate 4*K products;
 35 accumulate the 4*K products with a 32-bit single precision floating-point data element corresponding to a row m of the M rows, and a column n of the N columns, of the third matrix, to generate a result 32-bit single precision floating-point data element; and
 store the result 32-bit single precision floating-point data element at the row m and the column n of the third matrix.
- 40
- 45
- 50
- 55




FIG. 1A




FIG. 1B




FIG. 2

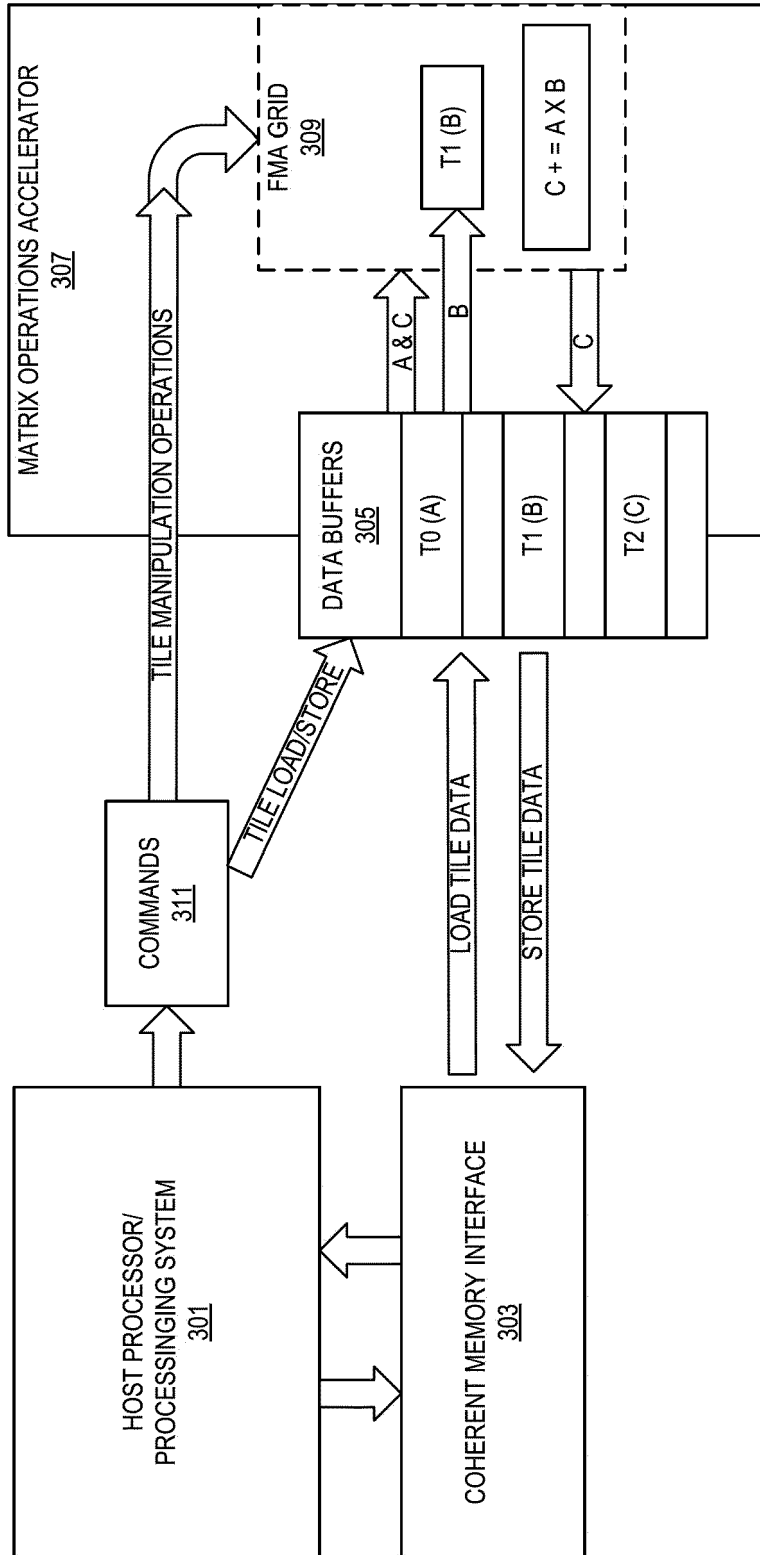


FIG. 3




FIG. 4




FIG. 5




FIG. 6

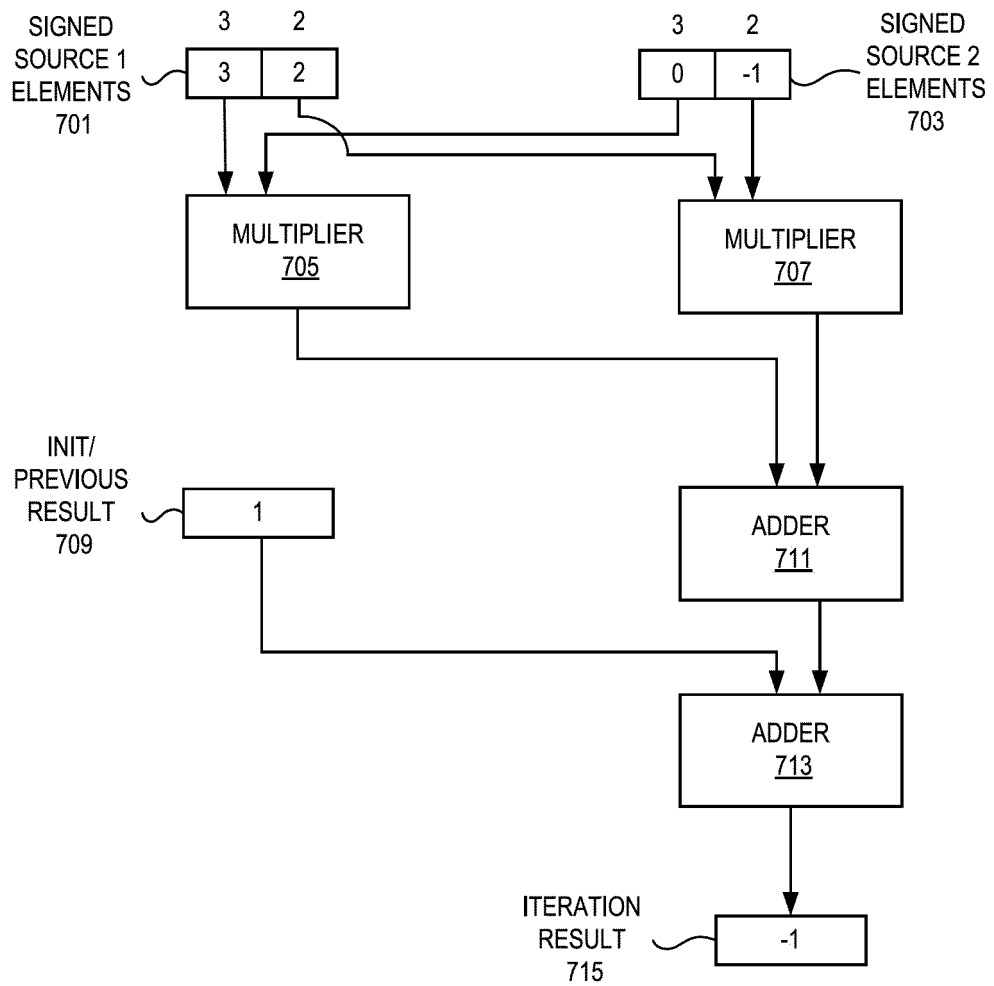


FIG. 7

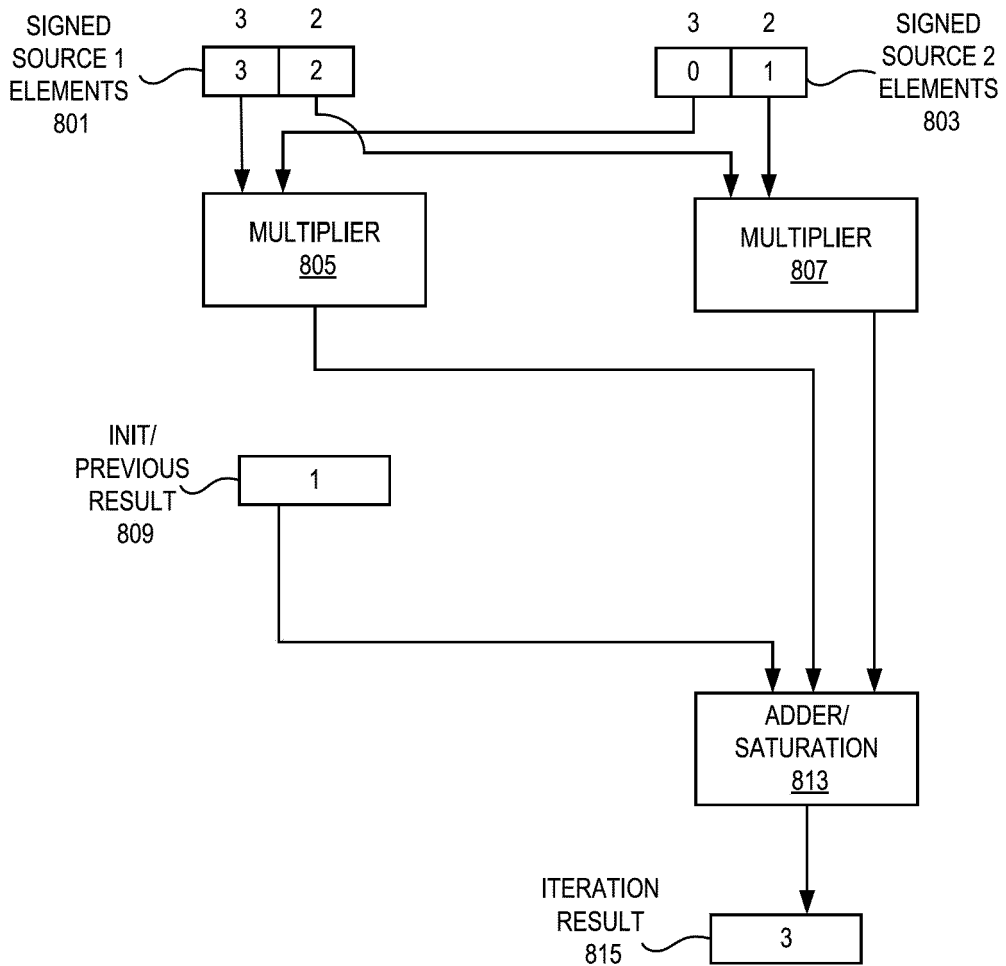


FIG. 8

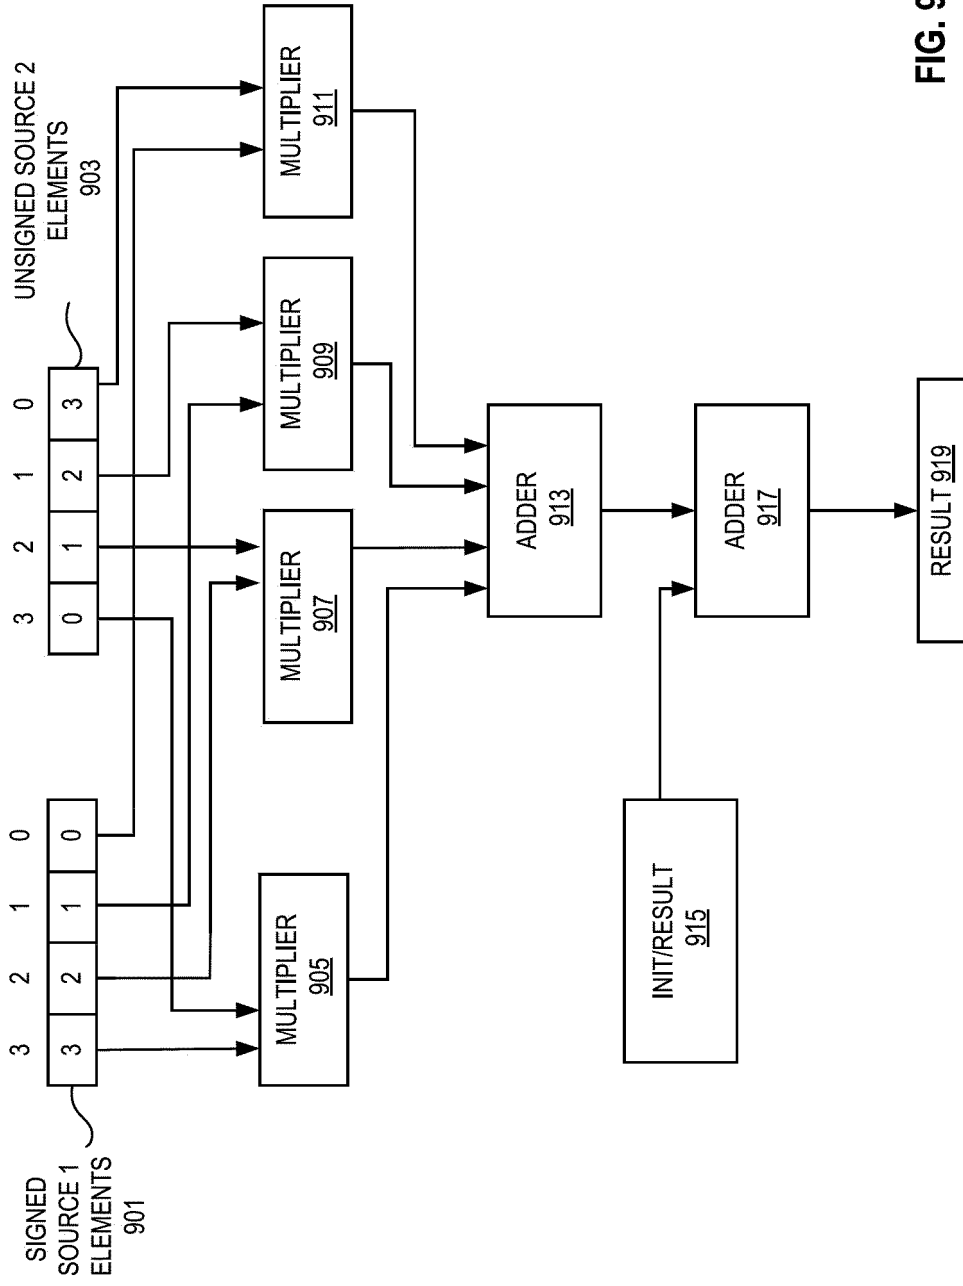


FIG. 9




FIG. 10

ACCUMULATOR 2X INPUT SIZES 1101

SOURCES	BITS	ACCUMULATOR	BITS
BYTE	8	WORD/HPFP	16
WORD	16	INT32/SPFP	32
SPFP/INT32	32	INT64/DPFP	64

ACCUMULATOR 4X INPUT SIZES 1103

SOURCES	BITS	ACCUMULATOR	BITS
BYTE	8	INT32/SPFP	32
WORD	16	INT64/DPFP	64

ACCUMULATOR 8X INPUT SIZES 1105

SOURCES	BITS	ACCUMULATOR	BITS
BYTE	8	INT64/DPFP	64

FIG. 11

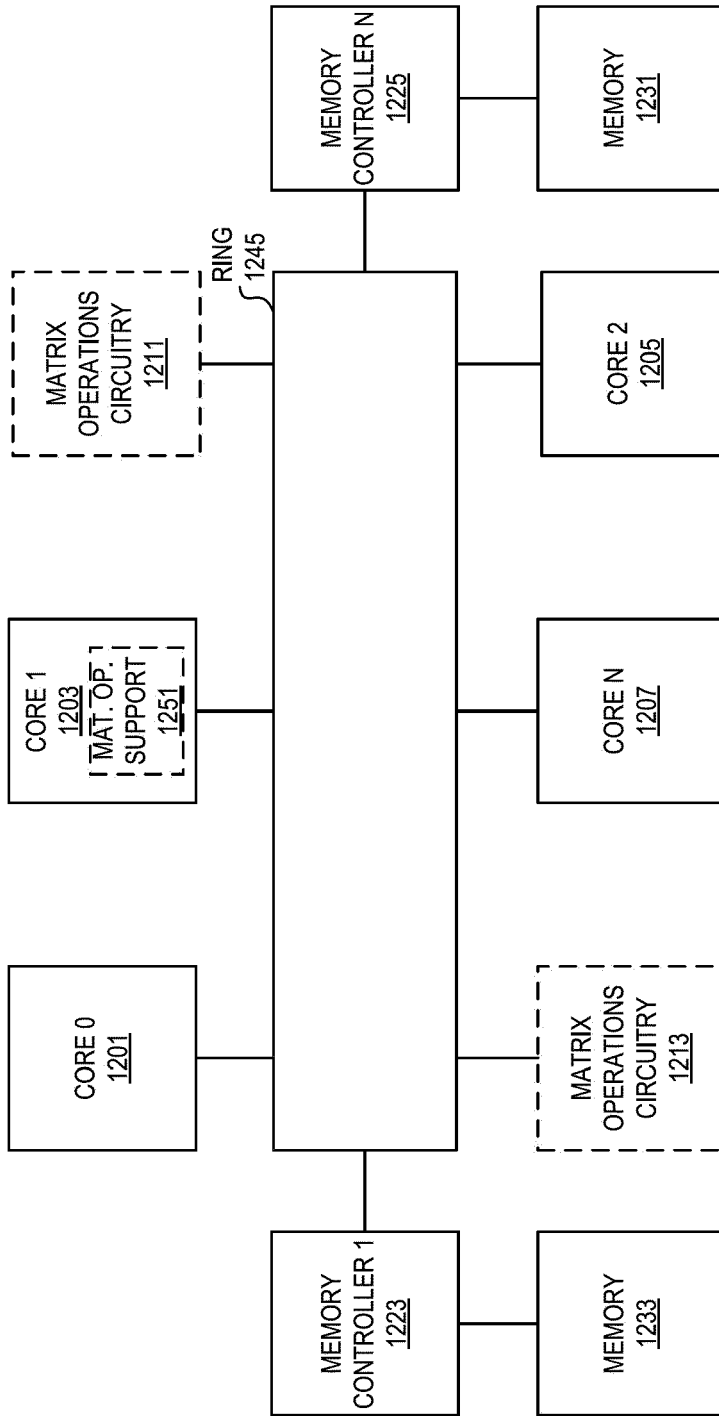


FIG. 12

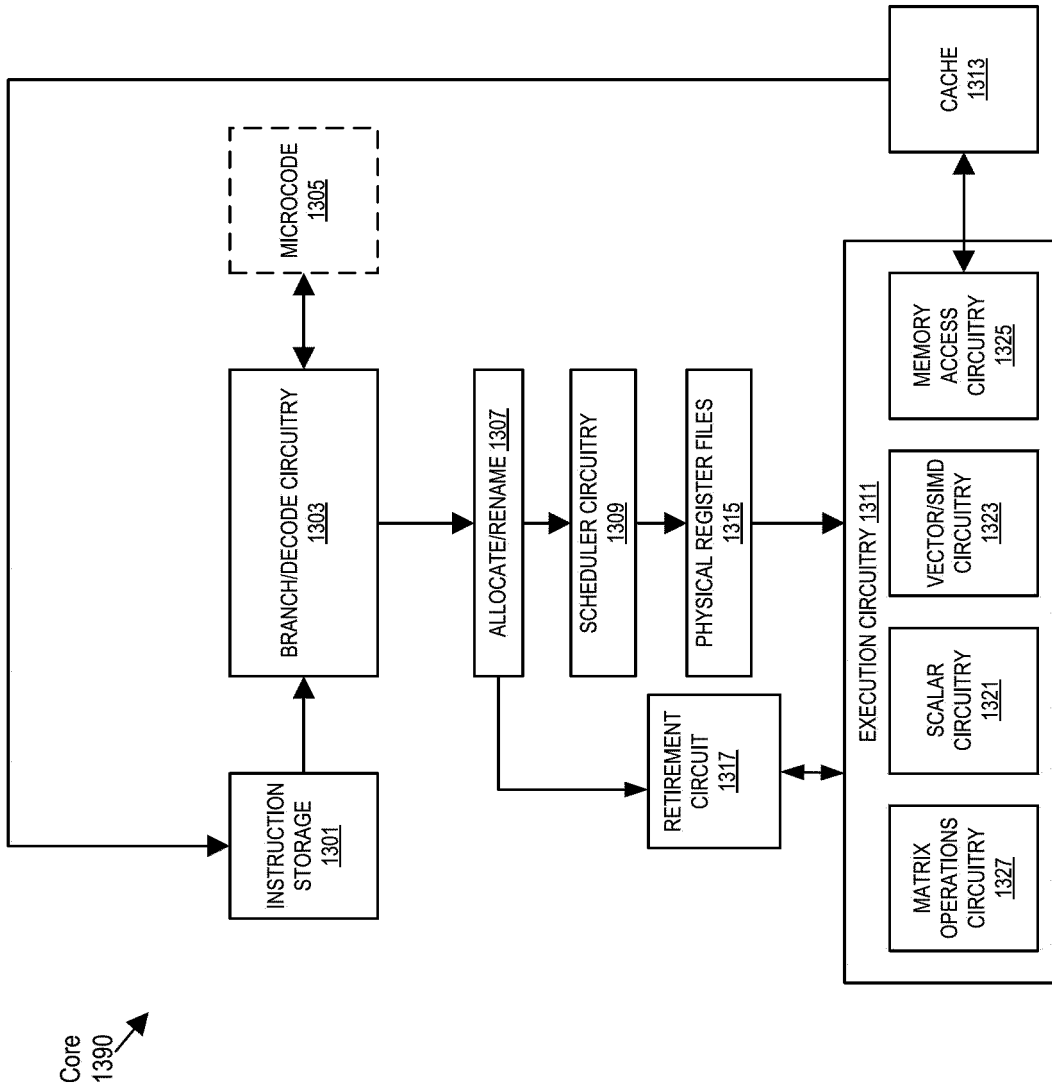


FIG. 13

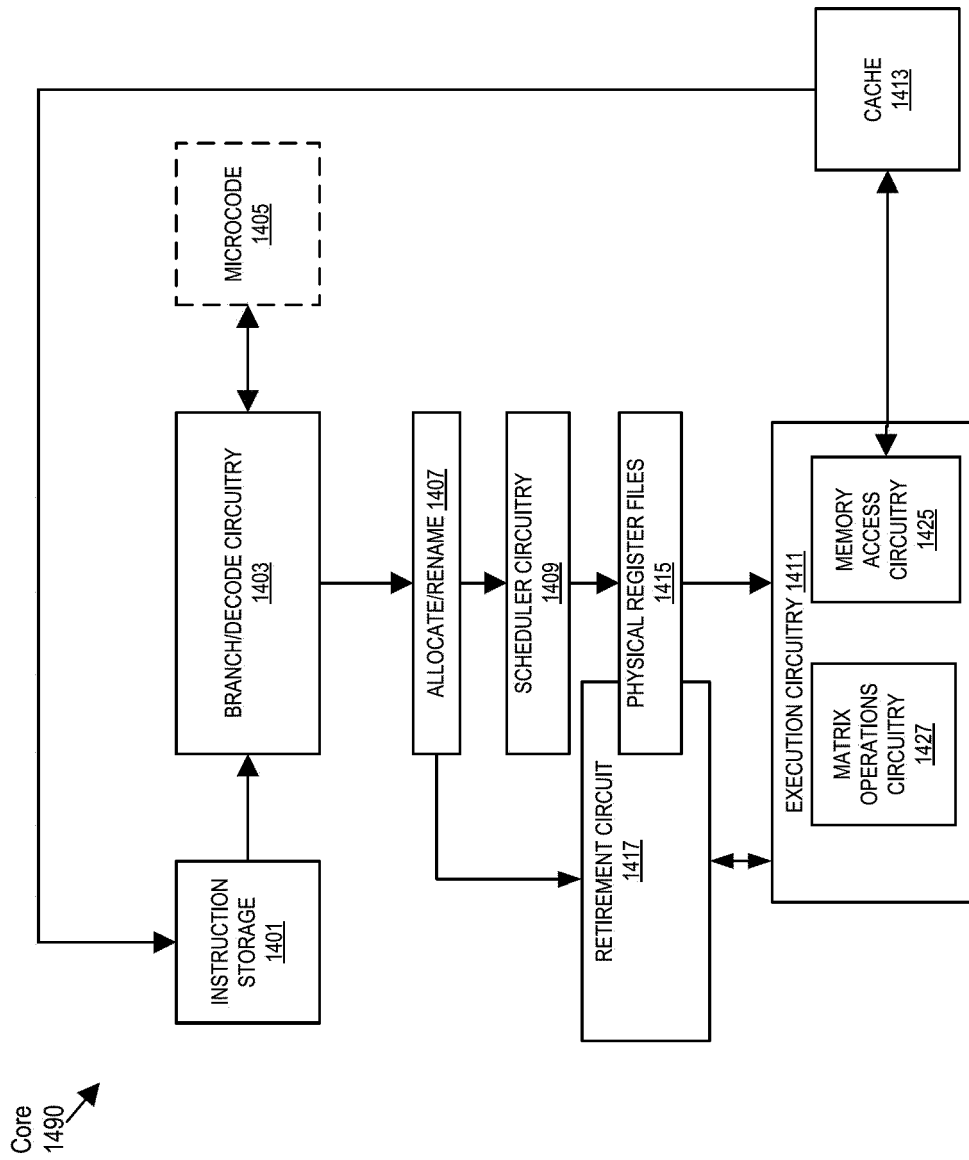


FIG. 14

$$A = \begin{bmatrix} A_{11} & A_{12} & A_{13} \\ A_{21} & A_{22} & A_{23} \end{bmatrix}$$


ADDR	VALUE
0	A_{11}
1	A_{12}
2	A_{13}
3	A_{21}
4	A_{22}
5	A_{23}

ROW MAJOR

ADDR	VALUE
0	A_{11}
1	A_{21}
2	A_{12}
3	A_{22}
4	A_{13}
5	A_{23}

COLUMN MAJOR

FIG. 15



```

TILECONFIG [RAX]
// ASSUME SOME OUTER LOOPS DRIVING THE CACHE TILING (NOT SHOWN)
{
  TILELOAD TMM0, RSI+RDI // SRC DST, RSI POINTS TO C, RDI HAS
  TILELOAD TMM1, RSI+RDI+N // SECOND TILE OF C, UNROLLING IN SIMD DIMENSION N
  MOV KK, 0
  LOOP:
  TILELOAD TMM2, R8+R9 // SRC2 IS STRIDED LOAD OF A, REUSED FOR 2 TMM A INSTR.
  TILELOAD TMM3, R10+R11 // SRC1 IS STRIDED LOAD OF B
  TMMAPS TMM0, TMM2, TMM2, TMM3 // UPDATE LEFT TILE OF C
  TILELOAD TMM3, R10+R11+N // SRC1 LOADED WITH B FROM NEXT RIGHTMOST TILE
  TMMAPS TMM1, TMM2, TMM2, TMM3 // UPDATE RIGHT TILE OF C
  ADD R8, K // UPDATE POINTERS BY CONSTANTS KNOWN OUTSIDE OF LOOP
  ADD R10, K*R11
  ADD KK, K
  CMP KK, LIMIT
  JNE LOOP
  TILESTORE RSI+RDI, TMM0 // UPDATE THE C MATRIX IN MEMORY
  TILESTORE RSI+RDI+M, TMM1
} // END OF OUTER LOOP
TILERELASE // RETURN TILES TO INIT STATE

```

FIG. 16

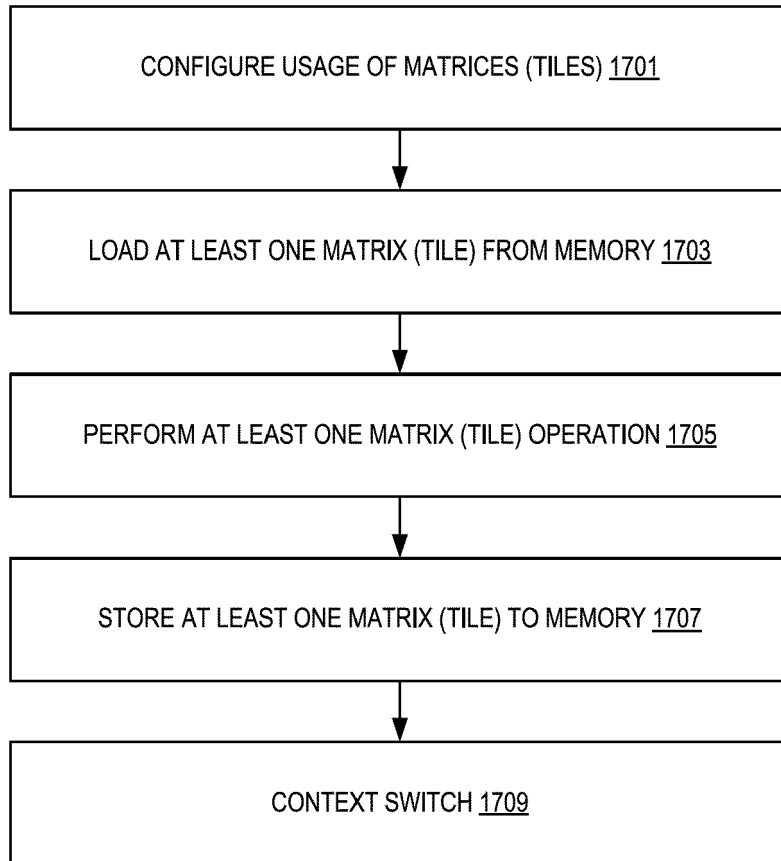


FIG. 17

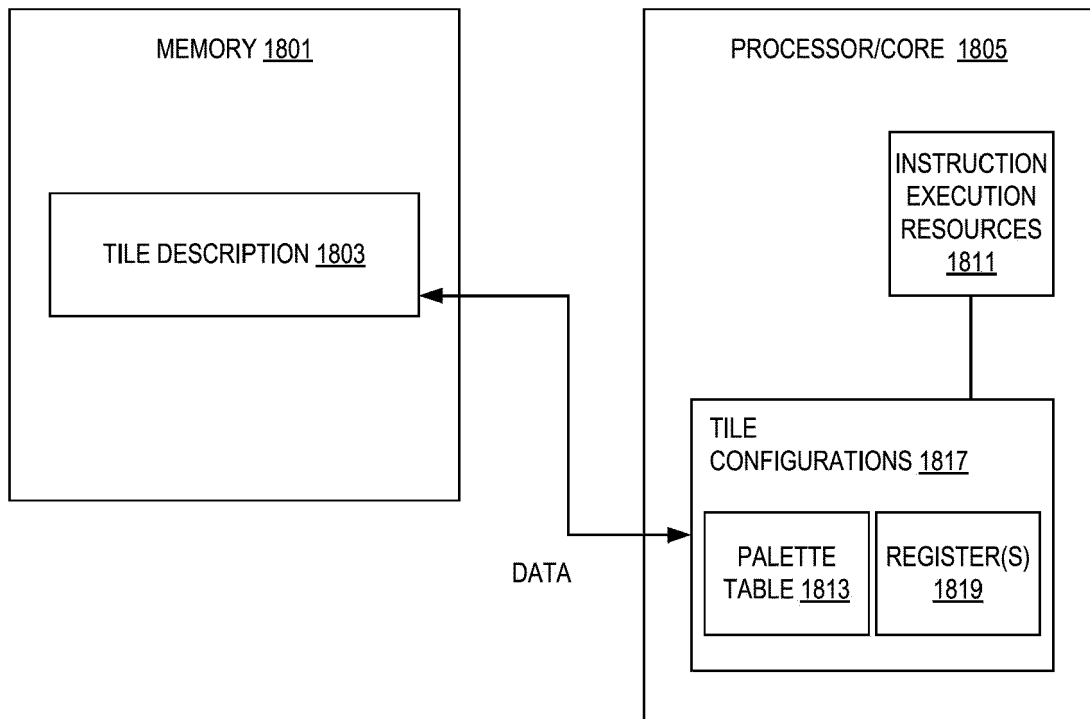


FIG. 18

PALETTE ID <u>1901</u>	STARTM <u>1903</u>
STARTP <u>1905</u>	PAIR INDICATORS <u>1907</u>
0	0
0	0

...

0	0
TMM0 ROWS <u>1913</u>	TMM0 COLUMNS <u>1915</u>
TMM1 ROWS	TMM1 COLUMNS
▪ ▪ ▪	
TMM15 ROWS	TMM15 COLUMNS
0	

FIG. 19

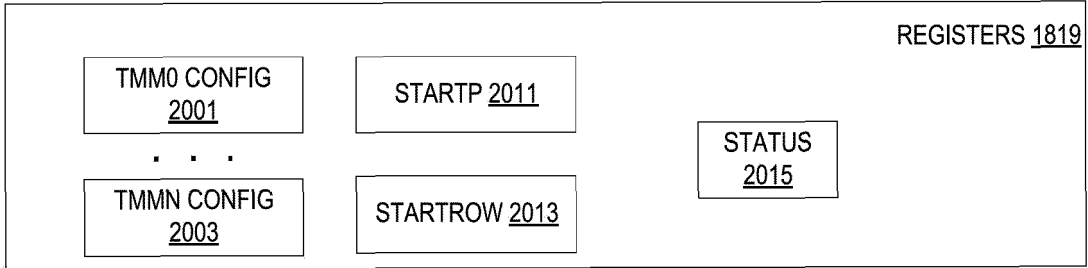


FIG. 20(A)

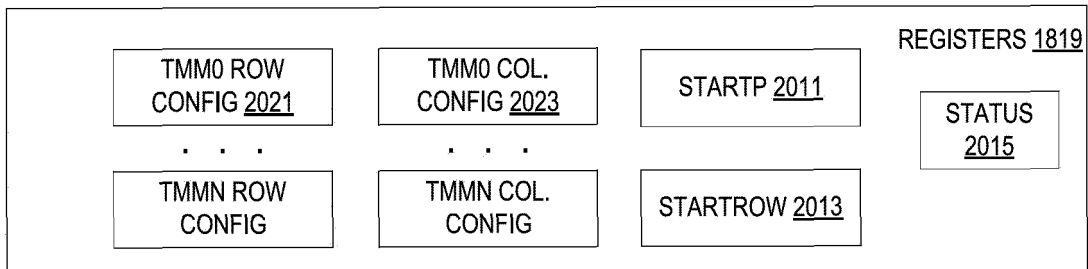


FIG. 20(B)




FIG. 20(C)

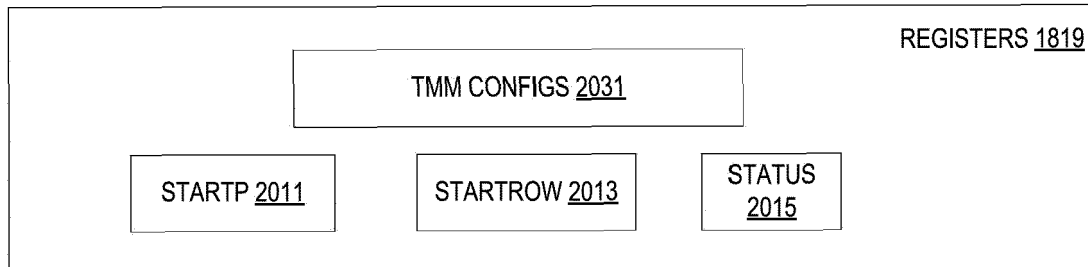


FIG. 20(D)

FIG. 211A

INSTRUCTION 2101

OPCODE TDPBF8PS* 2102	DESTINATION LOCATION 2104	FIRST SOURCE LOCATION 2106	SECOND SOURCE LOCATION 2108
-----------------------------	---------------------------------	----------------------------------	-----------------------------------

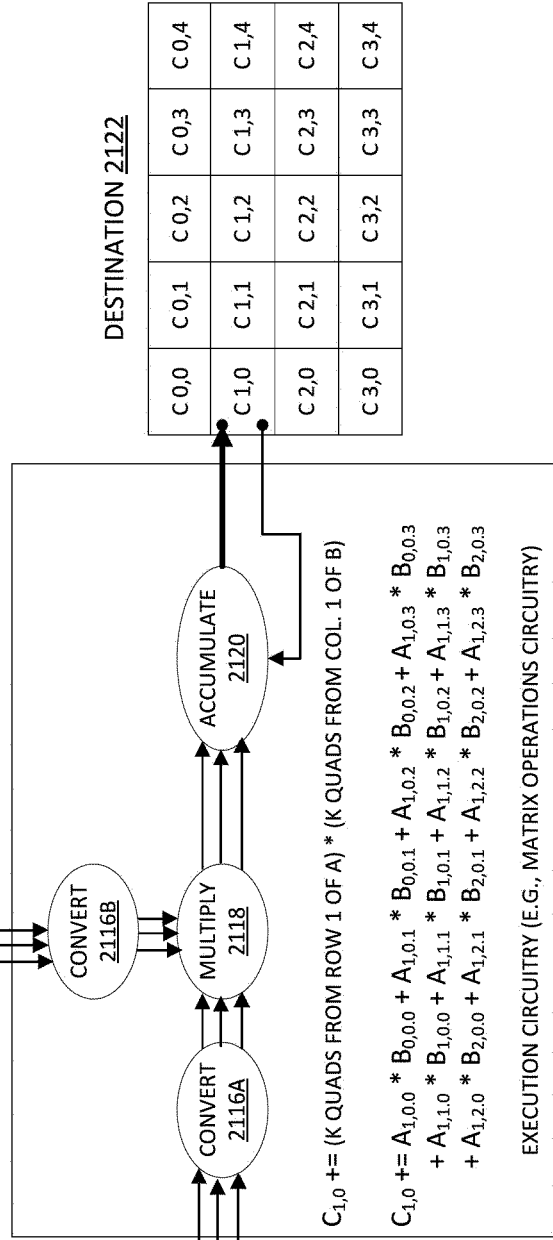
2100

B 0,0 0,1,2,3	B 0,1 0,1,2,3	B 0,2 0,1,2,3	B 0,3 0,1,2,3	B 0,4 0,1,2,3
B 1,0 0,1,2,3	B 1,1 0,1,2,3	B 1,2 0,1,2,3	B 1,3 0,1,2,3	B 1,4 0,1,2,3
B 2,0 0,1,2,3	B 2,1 0,1,2,3	B 2,2 0,1,2,3	B 2,3 0,1,2,3	B 2,4 0,1,2,3

SECOND
SOURCE
2112B

FIRST
SOURCE
2112A

A 0,0 0,1,2,3	A 0,1 0,1,2,3	A 0,2 0,1,2,3
A 1,0 0,1,2,3	A 1,1 0,1,2,3	A 1,2 0,1,2,3
A 2,0 0,1,2,3	A 2,1 0,1,2,3	A 2,2 0,1,2,3
A 3,0 0,1,2,3	A 3,1 0,1,2,3	A 3,2 0,1,2,3



DESTINATION 2122

C 0,0	C 0,1	C 0,2	C 0,3	C 0,4
C 1,0	C 1,1	C 1,2	C 1,3	C 1,4
C 2,0	C 2,1	C 2,2	C 2,3	C 2,4
C 3,0	C 3,1	C 3,2	C 3,3	C 3,4

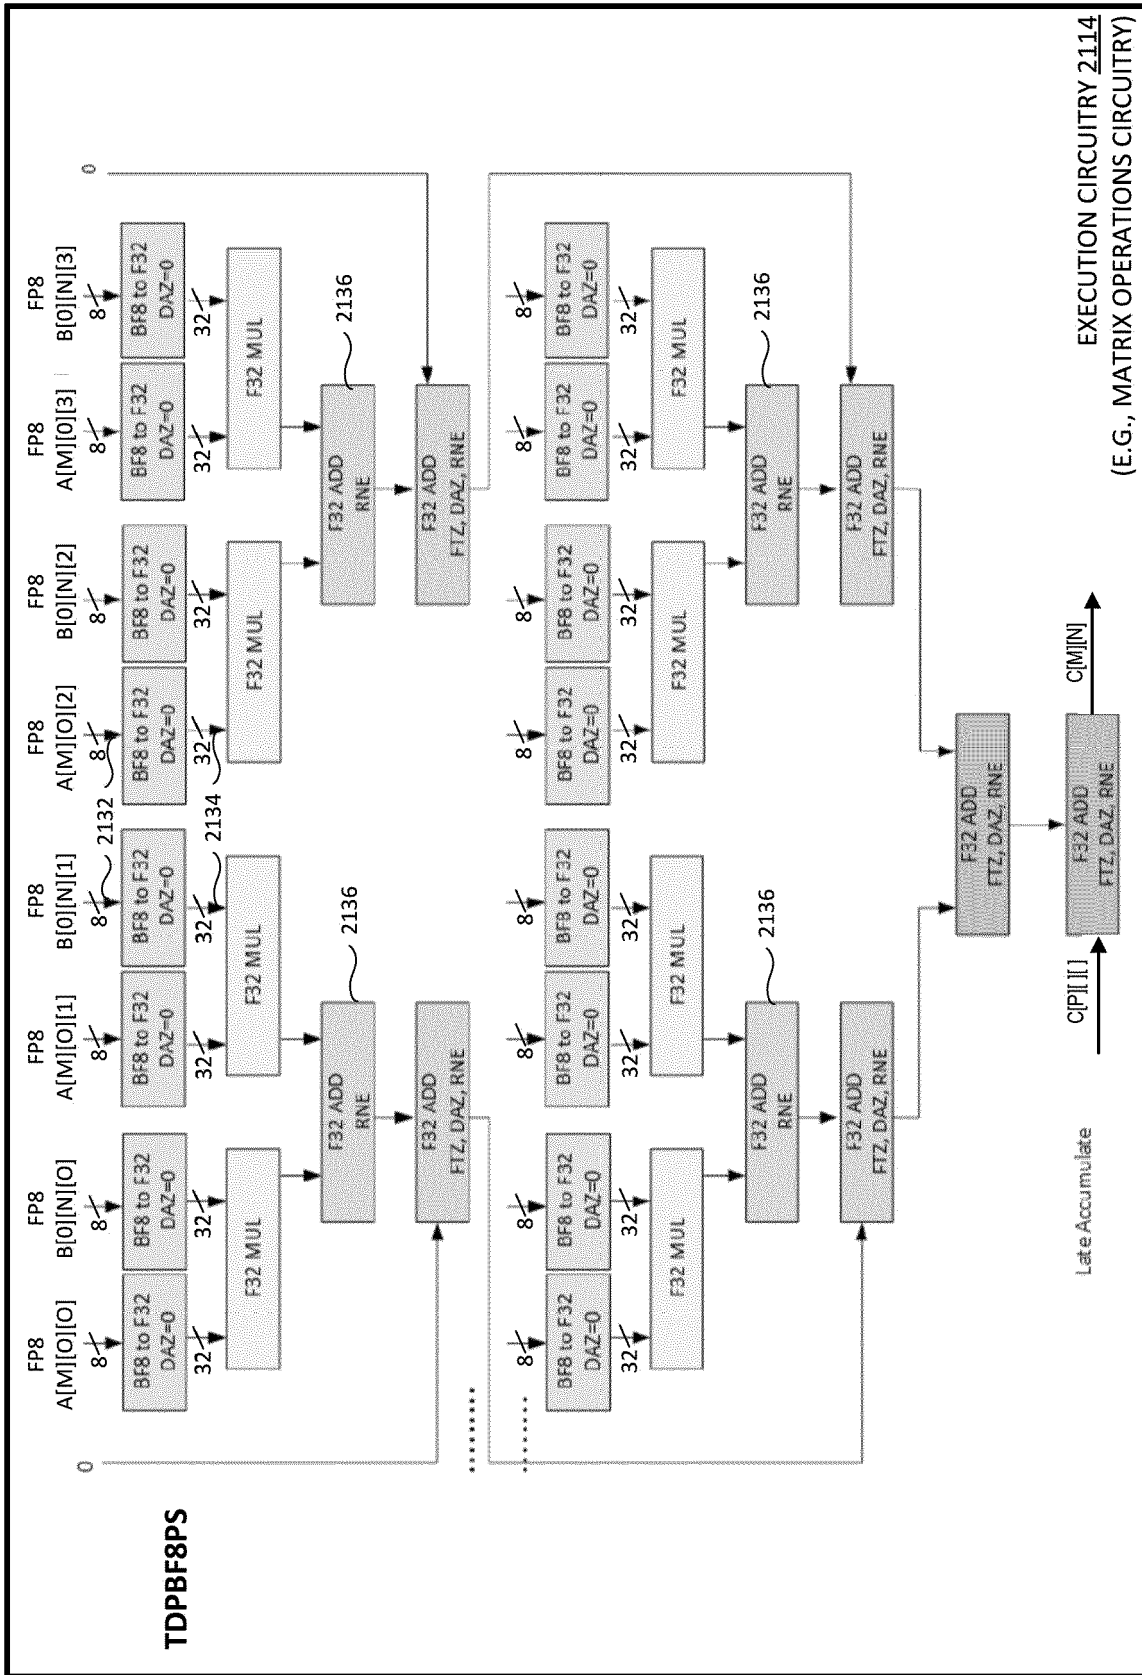


FIG. 21B

INSTRUCTION 2201

FIG. 22A

OPCODE TDPBF8PS* <u>2202</u>	DESTINATION LOCATION <u>2204</u>	FIRST SOURCE LOCATION <u>2206</u>	SECOND SOURCE LOCATION <u>2208</u>
------------------------------------	--	---	--

PSEUDOCODE 2200

```

TDPBF8PS tsrctest, tsrc1, tsrc2
// C = m x n (tsrctest), A = m x k (tsrc1), B = k x n (tsrc2)

# src1 and src2 elements are 4-tuples of bfloat8
elements_src1 := tsrc1.colsb / 4
elements_src2 := tsrc2.colsb / 4
elements_dest := tsrctest.colsb / 4
elements_temp1 := tsrctest.colsb / 2
elements_temp2 := tsrctest.colsb
elements_temp3 := tsrctest.colsb / 2

for m in 0 ... tsrctest.rows-1:
  temp1[ 0 ... elements_temp1-1 ] := 0
  for k in 0 ... elements_src1-1:
    for n in 0 ... elements_dest-1:

      // FP32 MUL with DAZ=1, FTZ=1, RNE rounding.
      // MXCSR is neither consulted nor updated.
      // No exceptions raised or denoted.
      temp2.fp32[4*n+0] = convert_bf8_to_fp32(tsrc1.row[m].bfloat8[4*k+0]) *
                          convert_bf8_to_fp32(tsrc2.row[k].bfloat8[4*n+0]);
      temp2.fp32[4*n+1] = convert_bf8_to_fp32(tsrc1.row[m].bfloat8[4*k+1]) *
                          convert_bf8_to_fp32(tsrc2.row[k].bfloat8[4*n+1]);
      temp2.fp32[4*n+2] = convert_bf8_to_fp32(tsrc1.row[m].bfloat8[4*k+2]) *
                          convert_bf8_to_fp32(tsrc2.row[k].bfloat8[4*n+2]);
      temp2.fp32[4*n+3] = convert_bf8_to_fp32(tsrc1.row[m].bfloat8[4*k+3]) *
                          convert_bf8_to_fp32(tsrc2.row[k].bfloat8[4*n+3]);

      // FP32 ADD, RNE rounding
      // MXCSR is neither consulted nor updated.
      // No exceptions raised or denoted.
      temp3.fp32[2*n+0] = temp2.fp32[4*n+0] + temp2.fp32[4*n+1];
      temp3.fp32[2*n+1] = temp2.fp32[4*n+2] + temp2.fp32[4*n+3];

      // FP32 ADD with DAZ=FTZ=1, RNE rounding
      // MXCSR is neither consulted nor updated.
      // No exceptions raised or denoted.
      temp1.fp32[2*n+0] += temp3.fp32[2*n+0];
      temp1.fp32[2*n+1] += temp3.fp32[2*n+1];

for n in 0 ... elements_dest-1:
  // DAZ=FTZ=1, RNE rounding.
  // MXCSR is neither consulted nor updated.
  // No exceptions raised or denoted.
  tmpf32 := temp1.fp32[2*n] + temp1.fp32[2*n+1]
  srcctest.row[m].fp32[n] := srcctest.row[m].fp32[n] + tmpf32

write_row_and_zero(tsrctest, m, tmp, tsrctest.colsb)
zero_upper_rows(tsrctest, tsrctest.rows)
zero_tileconfig_start()

```

PSEUDOCODE FOR HELPER FUNCTIONS 2200

```

define convert_bf8_to_fp32(x):
// The x parameter is bfloat8. Pack it in to upper 8b of a word.
// and then convert to fp32 while making fp16 denormals normal in FP32
// So DAZ = 0 for this operation
    word := 0
    word[15:8] := x
return cvt_fp16_fp32(word);

define write_row_and_zero(treg, r, data, nbytes):
    for j in 0 ... nbytes-1:
        treg.row[r].byte[j] := data.byte[j]
    // zero the rest of the row
    for j in nbytes ... palette_table[tileconfig.palette_id].bytes_per_row-1:
        treg.row[r].byte[j] := 0

define zero_upper_rows(treg, r):
    for i in r ... palette_table[tileconfig.palette_id].max_rows-1:
        for j in 0 ... palette_table[tileconfig.palette_id].bytes_per_row-1:
            treg.row[i].byte[j] := 0

Define zero_tileconfig_start():
    tileconfig.startRow := 0

```

FIG. 22B

2300

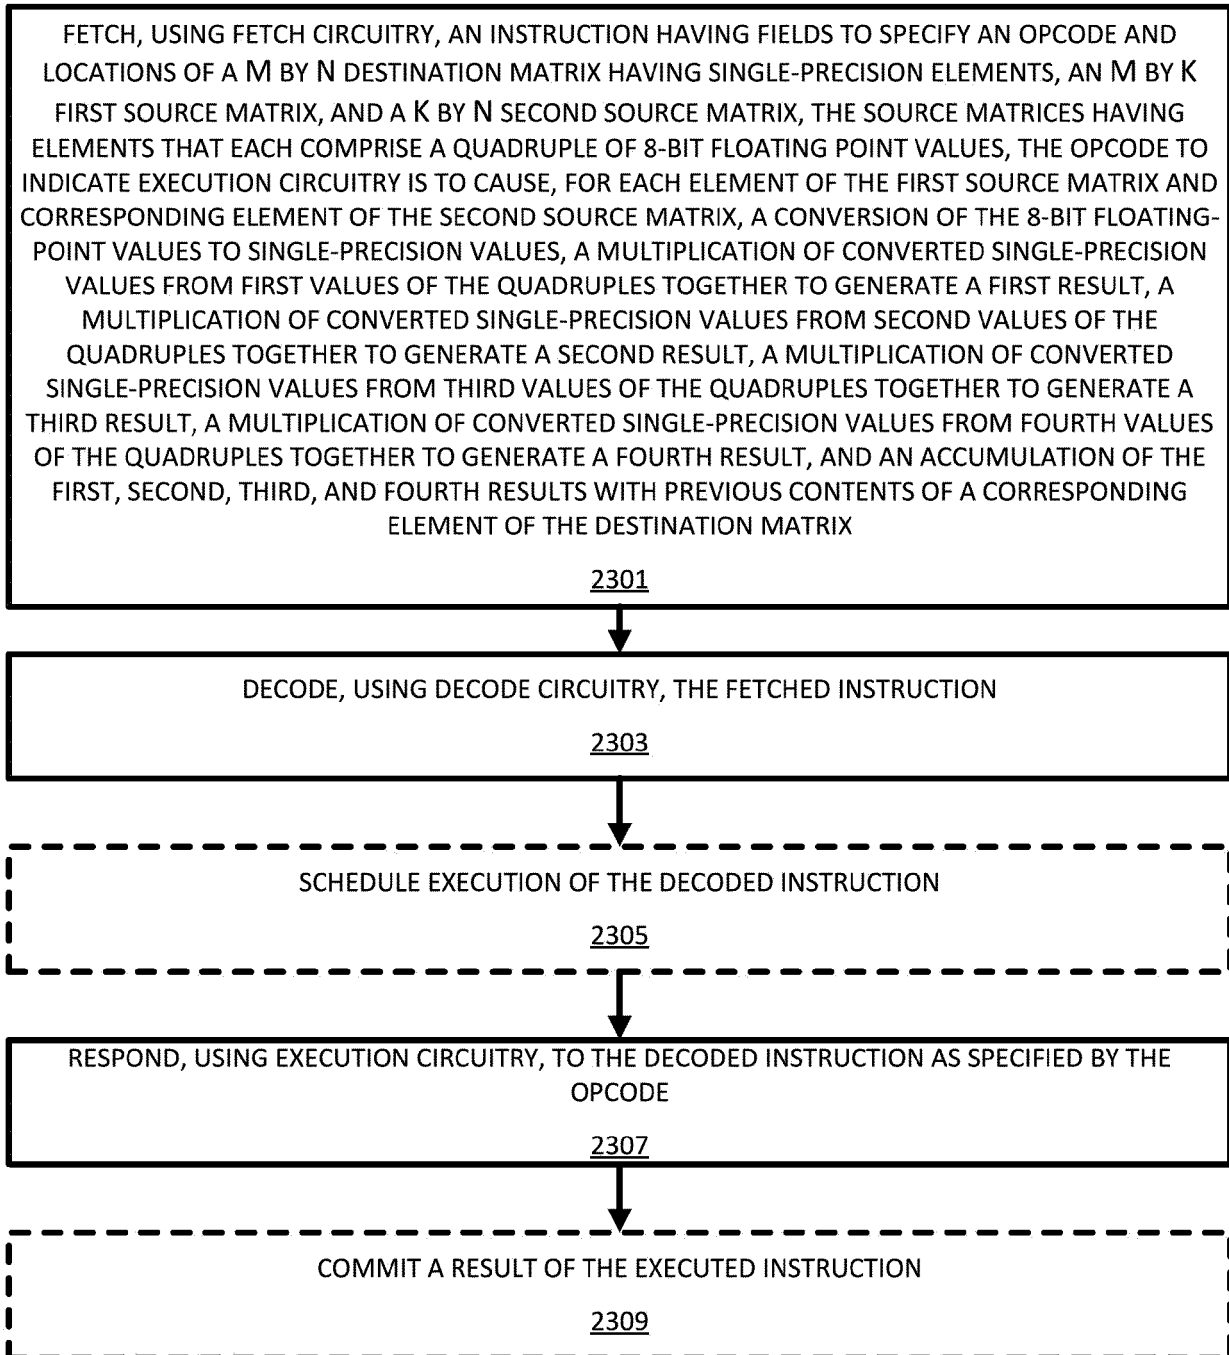


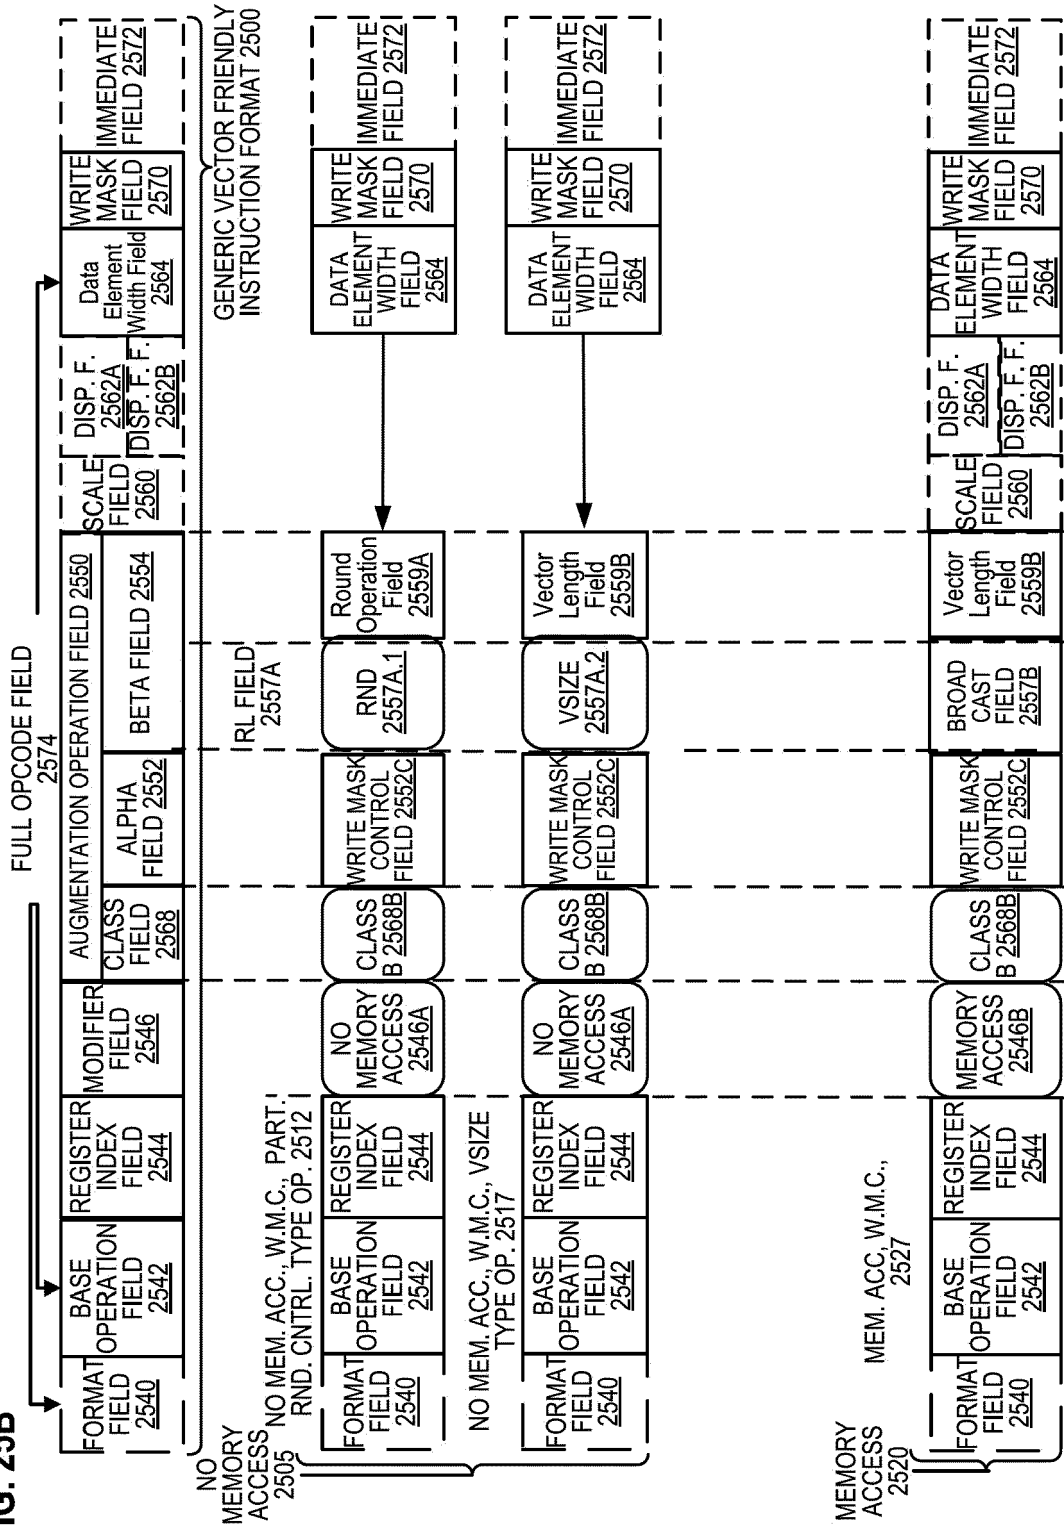
FIG. 23

TILE 8-BIT DOT PRODUCT (TDPBF8PS) INSTRUCTION 2400

OPCODE TDPBF8PS * <u>2402</u>	DESTINATION LOCATION <u>2404</u>	FIRST SOURCE LOCATION <u>2406</u>	SECOND SOURCE LOCATION <u>2408</u>	SOURCE ELEMENT FORMAT <u>2410</u>	(K)/(Z) <u>2412</u>	M (SOURCE 1 ROWS) <u>2414</u>	N (SOURCE 2 COLUMNS) <u>2416</u>
-------------------------------------	--	---	---	--	------------------------	--	---

FIG. 24

FIG. 25B






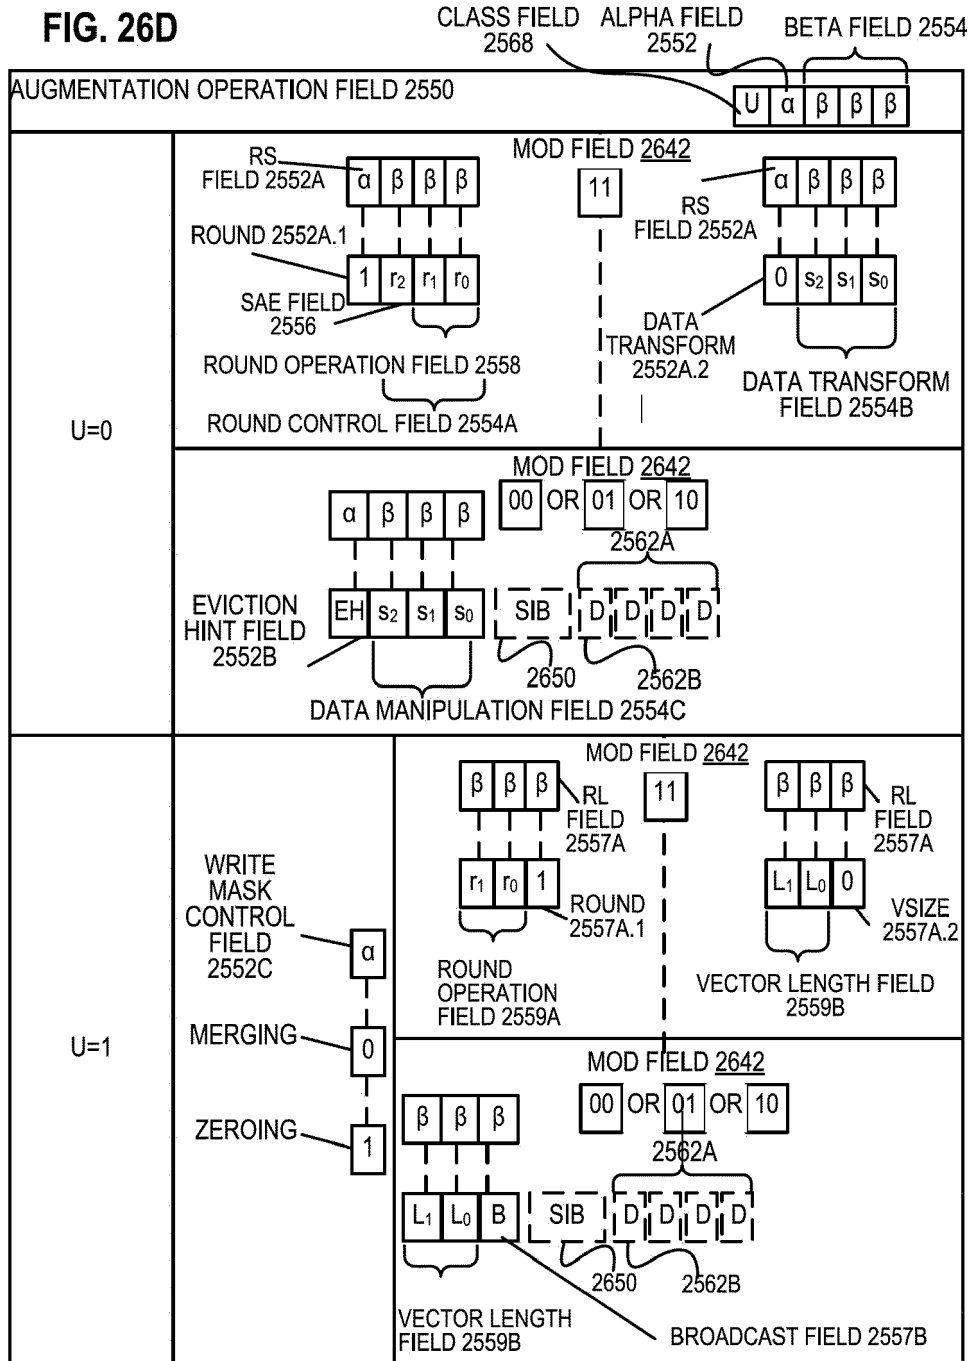


FIG. 26D






FIG. 27

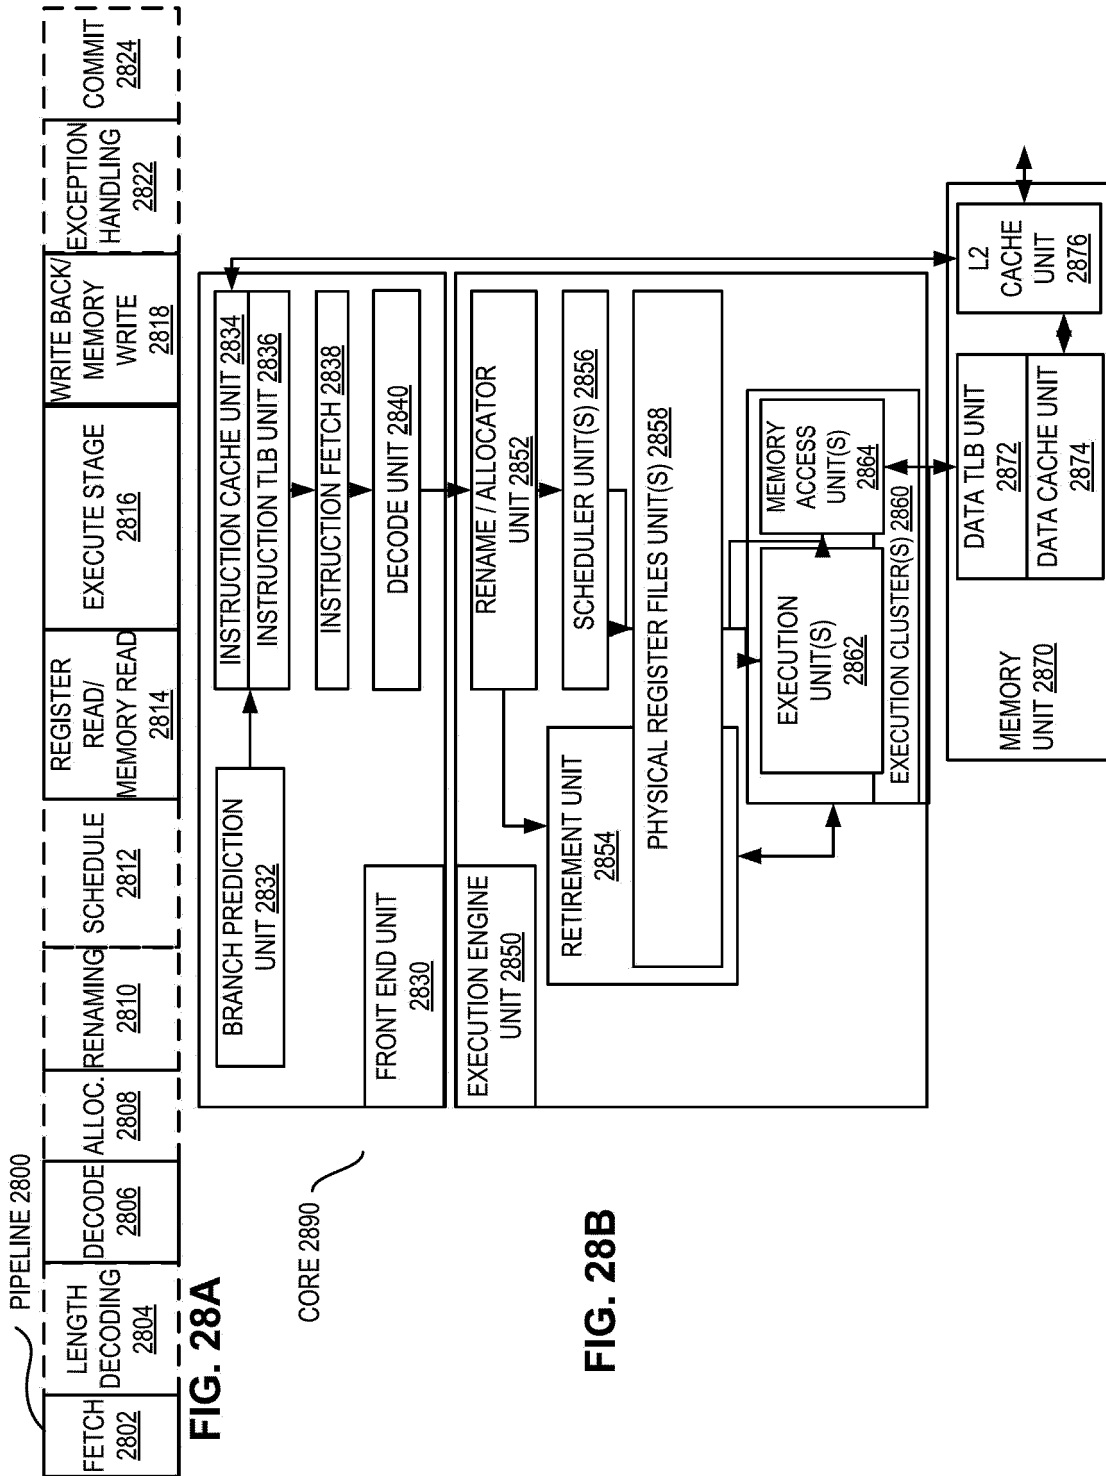


FIG. 28B

FIG. 29A





FIG. 29B






FIG. 30




FIG. 31




FIG. 32




FIG. 33




FIG. 34




FIG. 35

REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Non-patent literature cited in the description

- IA-32 Architectures Software Developer's Manual. September 2014 [0157]
- Intel® Advanced Vector Extensions Programming Reference, October 2014 [0157]